

On the Limit Imbalanced Logistic Regression by Binary Predictors

Vincent Runge^{*1}

¹B2PHI laboratory UMR 1181, INSERM, UVSQ, Institut Pasteur
Villejuif 94807, France.

Abstract

In this work, we introduce a modified (rescaled) likelihood for imbalanced logistic regression. This new approach makes easier the use of exponential priors and the computation of lasso regularization path. Precisely, we study a limiting behavior for which class imbalance is artificially increased by replication of the majority class observations. If some strong overlap conditions are satisfied, the maximum likelihood estimate converges towards a finite value close to the initial one (intercept excluded) as shown by simulations with binary predictors. This solution corresponds to the extremum of a concave function that we refer to as "rescaled" likelihood. In this context, the use of exponential priors has a clear interpretation as a shift on the predictor means for the minority class. Thanks to the simple binary structure, some random designs give analytic path estimators for the lasso regularization problem. An effective approximate path algorithm by piecewise logarithmic functions based on matrix inversions is also presented. This work was motivated by its potential application to spontaneous reports databases in a pharmacovigilance context.

Keywords: path estimator, pharamacovigilance model, piecewise logarithmic approximate path, limit class imbalance, rescaled likelihood, spontaneous reports database, square exact solution.

MS classification : Primary 62J12, 62F12, 62F15; secondary 34E05, 49M29, 62P10.

1 Introduction

If the response $y = 1$ is very rare compared with the response $y = 0$, we are in presence of a rare event configuration also called class imbalance. This problem recently got computer scientists' attention: they aimed at reducing computational costs by bypassing the class imbalance with resampling methods [12] [21] [6]. With these methods, the variance in estimating model parameters increases. Statisticians are aware of this problem and complex procedures such as local case-control sampling were proposed [8] (a method initiated in epidemiology [19]).

In a recent work (2007) by Art B. Owen [22], the opposite approach is considered: the class imbalance is infinitely increased in order to reach the theoretical distribution of the majority class observations. Owen proved that under some overlap conditions the model parameters are finite (apart from the intercept) and built a limit system of equations related to exponential tilting, whose solution is the new estimate. The resulting equations

^{*}E-mail: runge.vincent@gmail.com

include the distribution of the infinite class expressed through integrals, which are not easy to infer. This may explain that this work was broadly ignored (The author found it when Sections 2 and 3 were already completed).

In our approach, the observations of the majority class are infinitely replicated and the Owen's limit distribution becomes the observed distribution. This situation is a kind of degenerate case between resampling (we repeat observations) and infinitely class imbalance (the observed distribution is chosen as the theoretical one). Unlike Owen's result, our limit normal equations can be interpreted as the first order conditions of a new likelihood.

The idea of this work comes from the analysis of highly imbalanced binary spontaneous reports databases. Such databases are gathered by many countries and institutions (FDA, MHRA, WHO,...). Imbalanced logistic regression with binary predictors gives maximum likelihood estimate (MLE) very close to its limit imbalanced counterpart. This result makes possible the study of lasso-type regularization problem and the development of effective algorithms to provide model selection.

So far, only disproportionality methods are routinely used [18] for spontaneous report databases: predictors are analysed one by one, leading to a great number of false positive signals [13]. Mathematical tools adjusted to binary data for regression are surprisingly barely developed by scientists (only boolean matrices have been studied by some authors [16]). This results in an inflation of empiric methods using lasso regularization in recent years (from [3] to [1]). This is a worrying trend because recommendations made by these experts shift towards more complicated experimental methods and time-consuming algorithms, not towards a deeper mathematical understanding. This work is motivated by the need to better analyse this kind of applied problem.

The paper contains three main sections in which we present the following results:

- In Section 2, we investigate the properties of the logistic normal equations with binary predictors. Simple existence and uniqueness conditions of Silvapulle's type are found and some exact solutions presented. An invariance property in presence of intercept links this particular solution (called "square solution") to the limit imbalanced problem. We then acquaint ourselves with the issue of variance inflation of the imbalanced problem by computing the Fisher information.
- In Section 3, we derive Owen-type equations with a first order term evaluating the convergence rate. For the limit system of equations, the existence and uniqueness of the solution is proved with a new method leading to the minimization of a Kullback-Leibler divergence under linear constraints. A rescaling procedure on the initial likelihood and the previously found divergence justify the introduction of a rescaled likelihood corresponding to our limit imbalanced logistic regression problem. In a Bayesian framework, the Jeffreys penalty does not significantly decrease the variance of the estimator but other more appropriate priors, such that exponential ones, could help to reduce it (chosen according to the situation). The closeness in simulation between limit estimates and classical estimates compels us to go one step further with the study of regularization paths, in particular if the model is known to be sparse.
- In Section 4, we look at a lasso regularization problem for the rescaled likelihood, which has a clear interpretation as a shift on the predictor means for the class of

interest. We succeed in finding some path estimators in a few particular cases (independence and orthogonal design). In presence of correlation, we present an effective path following algorithm by piecewise logarithmic functions giving precise estimates. We conclude by explaining the need of an analysis of the correlation structure between predictors. This leads to simple algorithmic procedures with small computational costs for which many different prior penalties could be easily tested. Two examples are given using the French spontaneous reports database.

The expressions "infinitely imbalance" and "limit imbalance" are considered as synonymous, although we recommend the use of the second one in our context due to the simple unique limit we impose and an analogy with hydrodynamic limits (in fluid dynamics) while the first expression is related to the underlying distribution introduced by Owen.

We conclude this article by discussing the many opportunities that arise with the introduction of a rescaled likelihood in a Bayesian context and of the path following algorithm by logarithmic functions.

2 The logistic regression by binary predictors

2.1 Logistic normal equations

The binary logistic regression (BLR) problem consists in the determination of coefficients $\hat{\beta}$ maximizing a smooth and concave likelihood function given by the relation

$$L(\beta|I_0, I_1, n^0, n^1) = \prod_{i=1}^{q_1} \left(\frac{e^{(I_1\beta)_i}}{1 + e^{(I_1\beta)_i}} \right)^{n_i^1} \prod_{i=1}^{q_0} \left(1 + e^{(I_0\beta)_i} \right)^{-n_i^0},$$

where $\beta = (\beta_i) \in \mathbb{R}^{p+1}$ is indexed from zero with β_0 corresponding to the intercept. Binary design matrices $I_1 \in \mathcal{M}_{q_1 \times (p+1)}(\mathbb{B})$ and $I_0 \in \mathcal{M}_{q_0 \times (p+1)}(\mathbb{B})$ with $\mathbb{B} = \{0, 1\}$ are of full rank: they aggregate the p binary predictors. Vectors of weights $n^0 = (n_1^0, \dots, n_{q_0}^0)^T \in (\mathbb{N}^*)^{q_0}$ and $n^1 = (n_1^1, \dots, n_{q_1}^1)^T \in (\mathbb{N}^*)^{q_1}$ save repetitions for distinct observations in response classes 0 and 1 separately. The binary structure favours repetitions in the sequence of observations, which justifies these notations. Moreover $(I_0\beta)_i$ is the i -th component of vector $I_0\beta \in \mathbb{R}^{q_0}$ (the same for $(I_1\beta)_i$).

We introduce other notations thereafter used within this article. The modulus of a vector denotes its l^1 norm, while the overline sign on lower cases stands for l^1 normalization. For example $|n^1| = \sum_{i=1}^{q_1} n_i^1$ and $\bar{n}_i^1 = n_i^1/|n^1|$ gives the vector \bar{n}^1 . A_i is the i -th row of the matrix A and its roman upper case equivalent A is the matrix A in which the first column filled by ones (associated to the intercept) was removed. We also need $N^1 = \mathbf{I}_1^T n^1 \in \mathbb{R}^p$. An important feature in our study is the predictor means vector \bar{N}^1 for class 1 obtained by the relation $\mathbf{I}_1^T \bar{n}^1 = \bar{N}^1$. For vectors of same size $u, v \in \mathbb{R}^q$, uv is the vector with components $u_k v_k$, $k \in \{1, \dots, q\}$. $\tilde{\beta}$ is the vector β without the intercept coefficient β_0 . From Subsection 3.2, the notations I and \mathbf{I} for matrices I_0 and \mathbf{I}_0 respectively are often used (and $q = q_0$).

For ease of calculation, we consider the opposite of the log-likelihood. If $I_1 = I_0 = \mathcal{I}$, we have $q_1 = q_0 = q$ and we can introduce vectors $n = n^1 + n^0$ and $\Delta n = n^1 - n^0$. In this

latter case, we write

$$l(\beta) = -\log(L(\beta)) = |n| \log 2 + \sum_{i=1}^q \left(-\Delta n_i \left(\frac{1}{2} (\mathcal{I}\beta)_i \right) + n_i \log \cosh \left(\frac{1}{2} (\mathcal{I}\beta)_i \right) \right),$$

and first order conditions are computed, differentiating l with respect to each β_j coefficient. We obtain

$$0 = \frac{\partial l(\beta)}{\partial \beta_j} = \sum_{i=1}^q \left(-\Delta n_i \left(\frac{1}{2} \mathcal{I}_{ij} \right) + \frac{1}{2} n_i \mathcal{I}_{ij} \tanh \left(\frac{1}{2} (\mathcal{I}\beta)_i \right) \right), \quad j \in \{0, \dots, p\},$$

or in matrix form

$$\mathcal{I}^T \Delta n = \mathcal{I}^T \left(n \tanh \left(\frac{1}{2} \mathcal{I}\beta \right) \right). \quad (2.1)$$

In a general framework with non-identical matrices I_0 and I_1 , we likewise derive

$$I_1^T n^1 - I_0^T n^0 = I_1^T \left(n^1 \tanh \left(\frac{1}{2} I_1 \beta \right) \right) + I_0^T \left(n^0 \tanh \left(\frac{1}{2} I_0 \beta \right) \right). \quad (2.2)$$

Remark 2.1. *These equations are usually presented with a logistic function but we chose another expression to highlight the link with existence and uniqueness conditions.*

2.2 Existence and uniqueness

Necessary and sufficient conditions to ensure existence and uniqueness of the MLE are well-known, they were established by Silvapulle in 1981 [27]. They consist in satisfying an overlap condition $S_1 \cap S_0 \neq \emptyset$ with

$$S_1 = \{I_1^T s_1 \mid s_1 \in (\mathbb{R}_+^*)^{q_1}\}, \quad S_0 = \{I_0^T s_0 \mid s_0 \in (\mathbb{R}_+^*)^{q_0}\}.$$

For the BLR problem, a more convenient description is possible:

Theorem 2.1. *The BLR problem admits a unique solution if and only if there exist $n^+ \in (\mathbb{N}^*)^{q_1}$ and $n^* \in (\mathbb{N}^*)^{q_0}$, such that $I_1^T n^+ = I_0^T n^*$.*

Looking at equations (2.2), this theorem means that a MLE exists and is unique if one can find a couple (n^+, n^*) of observations of the rows in I_1 and in I_0 such that $|n^+| = |n^*|$ vanishing all the regression coefficients (intercept included). An easy necessary condition to check is that at least one 0 and one 1 are present in each column of I_0 and I_1 .

Proof. If $I_1^T n^+ = I_0^T n^*$, the Silvapulle's condition is immediately verified. Reciprocally, $S_1 \cap S_0$ is an open subset of \mathbb{R}^{p+1} with positive measure because I_0 and I_1 are full rank matrices. By a density argument, there exist $q \in (\mathbb{Q} \cap]0, 1])^{p+1}$, $\lambda \in (\mathbb{R}_+^*)^{q_0}$ and $\mu \in (\mathbb{R}_+^*)^{q_1}$ satisfying $I_0^T \lambda = I_1^T \mu = q$. We reorder the rows in I_0 and I_1 such that the first $p+1$ rows are linearly independent. Let H_0 in $\mathcal{M}_{q_0 \times (q_0 - p - 1)}(\mathbb{R})$ and H_1 in $\mathcal{M}_{q_1 \times (q_1 - p - 1)}(\mathbb{R})$ be orthogonal matrices to I_0 and I_1 respectively. Because of the reorganization of the rows in I_i ($i \in \{0, 1\}$) we can choose a H_i where its last $q_i - p - 1$ rows form an identity matrix $\mathbb{I}_{q_i - p - 1}$. For all $\alpha_0 \in \mathbb{R}^{q_0 - p - 1}$ and $\alpha_1 \in \mathbb{R}^{q_1 - p - 1}$ we have the relation $I_0^T (\lambda + H_0 \alpha_0) = I_1^T (\mu + H_1 \alpha_1) = q$. Again with a density argument, we find α_0 such that $\lambda_i + (\alpha_0)_i \in \mathbb{Q}_+^*$ for all $i \in \{p+2, \dots, q_0\}$ and satisfying the constraint $\lambda + H_0 \alpha_0 \in (\mathbb{R}_+^*)^{q_0}$. For a matrix $A \in \mathcal{M}_{n \times m}(\mathbb{R})$, a vector $v \in \mathbb{R}^m$ and $J \subset \{1, \dots, m\}$, let $[Av]_J$ denote the vector $A_J v_J$, where A_J (resp. v_J) corresponds to the submatrix of A (resp. subvector of v) obtained by removing from A (resp. from v) the columns (resp. rows) that do not

correspond to the indices in J . With this notation, we have $[I_0^T(\lambda + H_0\alpha_0)]_{\{1, \dots, p+1\}} = q - [I_0^T(\lambda + H_0\alpha_0)]_{\{p+2, \dots, q_0\}} \in \mathbb{Q}^{p+1}$. The binary matrix $(I_0^T)_{\{1, \dots, p+1\}}$ is then nonsingular and using its inverse in $\mathcal{M}_{(p+1) \times (p+1)}(\mathbb{Q})$ we obtain $(\lambda + H_0\alpha_0)_{\{1, \dots, p+1\}} \in (\mathbb{Q}_+^*)^{p+1}$. Finally $\alpha_0^* = \lambda + H_0\alpha_0 \in (\mathbb{Q}_+^*)^{q_0}$. The same arguments lead to a set of coefficients $\alpha_1^* = \mu + H_1\alpha_1 \in (\mathbb{Q}_+^*)^{q_1}$. Multiplying the vector (α_0^*, α_1^*) by the ppm of all its denominators proves the result. \square

2.3 The square case

The situation with identical square design matrices I_0 and I_1 leads to an analytic description of the MLE.

Theorem 2.2. *If $I_0 = I_1 = I$ is a square matrix, we have the following closed form for the maximum likelihood estimator:*

$$\hat{\beta} = I^{-1} \log \left(\frac{n^1}{n^0} \right). \quad (2.3)$$

Proof. the matrix I verifies the condition $q = p + 1$ and is nonsingular with I^{-1} its inverse (because I is of full rank). The vector T is defined as $T = \tanh\left(\frac{1}{2}I\beta\right) \in \mathbb{R}^{p+1}$ i.e. $I\beta = \log\left(\frac{1+T}{1-T}\right)$. Multiplying (2.1) by $(I^{-1})^T = (I^T)^{-1}$, we get $T = \frac{\Delta n}{n}$. Hence, $\beta = I^{-1} \log\left(\frac{n+\Delta n}{n-\Delta n}\right)$, which achieves the proof. \square

Remark 2.2. *If one of the components in the vectors of weights n^1 or n^0 vanishes, some of the regression coefficients become infinite (but not necessarily all of them).*

To our knowledge, this is the first general closed form found in the resolution of a logistic regression. There exist partial results for a unique categorical predictor exposed by Lipovetsky in 2014 [17]. An explanation for the lack of such a simple result stands in the poorly studied finite observation structure made possible through binary predictors. In Appendix A, some particular solutions to equations (2.3) are presented.

2.3.1 Invariance if intercept

We establish an invariance property making a link with the imbalanced problem.

Proposition 2.1. *In the square case with intercept, multiplying all the components of n^1 or n^0 by a same integer does not change the value of the MLE apart from the intercept.*

Proof. The inverse of a matrix with an intercept term verifies the relation

$$I^{-1} \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

which means that we can rewrite equations (2.3) as

$$\hat{\beta}_i = \log \left(\prod_{j=0}^p \left(\frac{n_j^1}{n_j^0} \right)^{a_{ij}} \right), \quad i \in \{0, \dots, p\}, \quad a_{ij} = (I^{-1})_{ij} \text{ and } \sum_{j=0}^p a_{ij} = \delta_{0i}.$$

Substituting n_j^0 by $s \times n_j^0$ (or n_j^1 by $s \times n_j^1$) with $s \in \mathbb{N}^*$ gives the same result for all β_i , $i \in \{1, \dots, p\}$. \square

2.3.2 Asymptotic variance

To conclude this section, we study the asymptotic behavior of the estimator for large $|n^1|$ and $|n^0|$. Since the MLE (intercept excluded) remains the same with or without a class imbalance (see the invariance property), we have a glimpse of a general property in class imbalance.

Proposition 2.2. *In the square case BLR problem, the variance of the maximum likelihood estimator is approximately given by relations*

$$V(\hat{\beta}_i) \approx \sum_{j=0}^p a_{ij}^2 \left(\frac{1}{n_j^1} + \frac{1}{n_j^0} \right), \quad i \in \{0, \dots, p\}, \quad a_{ij} = (I^{-1})_{ij}.$$

Proof. We compute the observed Fisher information $\mathcal{I}(\hat{\beta}) = I^T D I$ with D a diagonal matrix with elements $n_i \hat{p}_i (1 - \hat{p}_i)$ and $\hat{p}_i = 1/(1 + e^{-(I\hat{\beta})_i})$. Its inverse gives the desired result, knowing that $n_i \hat{p}_i = n_i^1$ and $n_i(1 - \hat{p}_i) = n_i^0$. \square

Another method uses the closed form (2.3) to perform variance and bias estimations by Taylor expansions with the multinomial random vector (n^1, n^0) . We obtain $V(\hat{\beta}_i) \approx \sum_{j=0}^p a_{ij}^2 (\frac{1}{n_j^1} + \frac{1}{n_j^0} - \frac{2}{|n|})$ and $Bias(\hat{\beta}_i) \approx \sum_{j=0}^p \frac{a_{ij}}{2} (\frac{1}{n_j^0} - \frac{1}{n_j^1})$, $i \in \{0, \dots, p\}$. However, simulations give inaccurate results and only the Fisher information method should be retained.

Given $|n|$ and $|\beta_0|$, we simulate 10^4 data sets with a different random binary square matrix I for each of them but a fixed model $(\beta_1, \dots, \beta_5) = (-0.5, -0.25, 0, 0.25, 0.5)$. We investigate the impact of the sample size and the value of the intercept term on the variance. In table 1, we compare the estimated standard deviation (*sd.*) with the Fisher standard deviation given in proposition 2.2 (*F.sd.*) accompanied by an estimation of the bias (*bias*) for coefficient $\beta_4 = 0.25$.

Table 1: Variance analysis in the square case

	β_0	-7	-6	-5	-4	-3	-2	-1	0
$ n $	$ n_0 / n_1 $	1052	385	142	52	19	7.1	2.7	1.0
10^3	sd.	0.768	0.467	0.333	0.286
	F.sd.	0.762	0.464	0.334	0.293
	bias	0.019	0.0037	0.0065	-0.0021
10^4	sd.	.	.	0.633	0.369	0.224	0.145	0.105	0.0931
	F.sd.	.	.	0.618	0.361	0.221	0.145	0.104	0.0918
	bias	.	.	0.019	9.4e-3	-2.4e-4	1.e-3	-3.0e-4	-2.9e-4
10^5	sd.	0.528	0.303	0.187	0.110	0.0690	0.0450	0.0328	0.0300
	F.sd.	0.512	0.301	0.183	0.111	0.0685	0.0451	0.0327	0.0291
	bias	0.020	5.4e-4	3.4e-3	-3.2e-5	-5.6e-5	4.5e-4	-2.8e-5	1.9e-4

These simulations highlight the accuracy of the "Fisher variance" in all configurations, which is very close to the estimated one. Bias is negligible compared with variance. Variance inflation is a key issue in class imbalance, we further explain how one can add a prior information to a rescaled likelihood to deal with this problem (see Subsection 3.4).

3 Limit imbalanced study

3.1 Owen-type equations

The limit case consists in infinitely replicating the majority class observations as if the theoretical distribution was the observed one. This is a degenerate case of the Owen's study, that is why we know that the intercept coefficient tends to minus infinity whereas other regression coefficients are finite if a stronger overlap condition is satisfied [22]. For the limit equations, an information reduction for the majority class occurs: only the means of the predictors matter, the correlation structure in this class of interest "disappears".

Proposition 3.1. *For an imbalanced binary logistic regression with a class size for response $y = 0$'s' times greater than the one for response $y = 1$, we obtained the following system of equations if a strong overlap condition is verified (see Theorem 3.2)*

$$\frac{\mathbf{n}_1^0}{n_1^0} = \bar{N}^1 + \frac{1}{s} \left(\frac{n_2^0}{(n_1^0)^2} \left(\frac{\mathbf{n}_2^0}{n_2^0} - \bar{N}^1 \right) - \frac{n_1^1}{n_1^0} \left(\frac{\mathbf{n}_1^1}{n_1^1} - \bar{N}^1 \right) \right) + o\left(\frac{1}{s}\right),$$

with $s = \frac{|n^0|}{|n^1|} \gg 1$. We used notations

$$n_1^0 = \sum_{i=1}^{q_0} \bar{n}_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}, \quad n_1^1 = \sum_{i=1}^{q_1} \bar{n}_i^1 e^{(\mathbf{I}_1 \tilde{\beta})_i}, \quad n_2^0 = \sum_{i=1}^{q_0} \bar{n}_i^0 e^{2(\mathbf{I}_0 \tilde{\beta})_i}, \quad n_2^1 = \sum_{i=1}^{q_1} \bar{n}_i^1 e^{2(\mathbf{I}_1 \tilde{\beta})_i},$$

and for vectors in \mathbb{R}^p :

$$\mathbf{n}_1^0 = \mathbf{I}_0^T (\bar{n}^0 e^{\mathbf{I}_0 \tilde{\beta}}), \quad \mathbf{n}_1^1 = \mathbf{I}_1^T (\bar{n}^1 e^{\mathbf{I}_1 \tilde{\beta}}), \quad \mathbf{n}_2^0 = \mathbf{I}_0^T (\bar{n}^0 e^{2\mathbf{I}_0 \tilde{\beta}}), \quad \mathbf{n}_2^1 = \mathbf{I}_1^T (\bar{n}^1 e^{2\mathbf{I}_1 \tilde{\beta}}).$$

The technical proof of this result is exposed in Appendix B.

As shown by simulations (see table 2), the first order and remainder terms are negligible quantities with binary predictors. This suggests the introduction of the following limit imbalanced equations.

Theorem 3.1. *For infinitely imbalanced binary logistic regression verifying a strong overlap condition (see Theorem 3.2), the following system of p limit imbalanced equations holds¹*

$$\mathbf{I}_0^T \left(\frac{\bar{n}^0 e^{\mathbf{I}_0 \tilde{\beta}}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \right) = \bar{N}^1. \quad (3.1)$$

Notice that the $\tilde{\beta}$ coefficients do not depend on the structure in rows of the design matrix associated to response $y = 1$ but only on the means of ones for each predictor: \bar{N}^1 .

Proof. For x near minus infinity, the hyperbolic tangent has the following first order expansion:

$$\tanh\left(\frac{x}{2}\right) = -1 + 2e^x + o(e^x).$$

From [22] we know that the intercept term tends to minus infinity, then with $x = \mathbf{I}_0 \beta$ or $x = \mathbf{I}_1 \beta$, we use the previous expansion neglecting the remainder term. Thus, equations (2.2) become

$$\mathbf{I}_1^T n^1 = \mathbf{I}_1^T \left(n^1 e^{\mathbf{I}_1 \beta} \right) + \mathbf{I}_0^T \left(n^0 e^{\mathbf{I}_0 \beta} \right), \quad (3.2)$$

¹With non-binary design matrices X_1 and X_0 and no vectors of weights, we obtain $\mathbf{x}_0^T \left(\frac{e^{\mathbf{x}_0 \tilde{\beta}}}{\sum_i e^{(\mathbf{x}_0 \tilde{\beta})_i}} \right) = \bar{N}^1$. These equations differ from Owen's [22].

and factoring by $\exp(\beta_0)$ in the first equation of this system we have

$$\exp(\beta_0) = \frac{|n^1|}{\sum_{i=1}^{q_1} n_i^1 e^{(\mathbf{I}_1 \tilde{\beta})_i} + \sum_{i=1}^{q_0} n_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \approx \frac{|n^1|}{\sum_{i=1}^{q_0} n_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}}, \quad (3.3)$$

because $\frac{|n^0|}{|n^1|} \rightarrow +\infty$. Looking back at (3.2) without the first equation, we have

$$\mathbf{I}_1^T \bar{n}^1 = \mathbf{I}_1^T \left(\frac{n^1 e^{\mathbf{I}_1 \tilde{\beta}}}{\sum_{i=1}^{q_0} n_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \right) + \mathbf{I}_0^T \left(\frac{n^0 e^{\mathbf{I}_0 \tilde{\beta}}}{\sum_{i=1}^{q_0} n_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \right)$$

but

$$\frac{n^1 e^{\mathbf{I}_1 \tilde{\beta}}}{\sum_{i=1}^{q_0} n_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} = \frac{|n^1|}{|n^0|} \frac{\bar{n}^1 e^{\mathbf{I}_1 \tilde{\beta}}}{\sum_{i=1}^{q_0} \bar{n}_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \rightarrow 0$$

because $\frac{|n^1|}{|n^0|} \rightarrow 0$ and we obtain the desired result. \square

In table 2, we present simulation results based on limit imbalanced equations (3.1) compared with classical logistic regression. The sample procedure is the same as the one used for table 1 except that we chose a fixed sample size $|n^1| + |n^0| = 10^4$ and a varying dimension for the matrix \mathbf{I}_0 ($q_0 = 10, 21, 32$).

Table 2: Variance and bias analysis in standard and imbalanced situations

q	β_0 $ n^0 / n^1 $	-5	-4	-3	-2	-1	0
		141	51	19	7.0	2.6	1.0
10	sd.	0.3614	0.2154	0.1331	0.08765	0.06384	0.05730
	sd. imb.	0.3614	0.2154	0.1331	0.08766	0.06402	0.05757
	bias	2.533e-3	7.572e-4	-9.553e-4	1.407e-3	2.797e-4	4.954e-4
	bias imb.	2.532e-3	7.527e-4	-9.433e-4	1.439e-3	3.261e-4	5.073e-4
	l^1	4.702e-4	6.849e-4	1.089e-3	1.762e-3	2.869e-3	4.806e-3
21	sd.	0.2643	0.1583	0.09879	0.06468	0.04839	0.04382
	sd. imb.	0.2643	0.1583	0.09879	0.06474	0.04860	0.04451
	bias	2.148e-3	1.956e-3	-1.583e-5	7.601e-4	2.461e-4	-6.280e-4
	bias imb.	2.130e-3	1.967e-3	-8.450e-6	7.752e-4	3.000e-4	-5.337e-4
	l^1	5.950e-4	8.864e-4	1.399e-3	2.281e-3	3.698e-3	5.952e-3
32	sd.	0.2438	0.1470	0.09346	0.06112	0.04585	0.04090
	sd. imb.	0.2438	0.1471	0.09349	0.06117	0.04628	0.04171
	bias	4.385e-3	1.034e-3	8.025e-4	3.584e-4	1.259e-3	-1.905e-4
	bias imb.	4.383e-3	1.048e-3	7.894e-4	3.537e-4	1.313e-3	-4.297e-5
	l^1	6.220e-4	9.239e-4	1.472e-3	2.341e-3	3.920e-3	6.306e-3

The two coefficients β_4 and $\tilde{\beta}_4$ are very close to each other as shown by the mean of the l^1 norm ($l^1 = \text{mean}|\beta_4 - \tilde{\beta}_4|$) – even if the problem is not imbalanced – so that standard deviation and bias are almost the same. This means that, if interesting properties can be established with the limit equations, this context will be appropriate to highlight new features in classical logistic regression.

The $1/s$ first order term in proposition 3.1 should be estimated to understand how good the limit imbalanced approximation is, without having to estimate the standard regression coefficients. Simulations show that this term is very small and we choose not to dwell on this intermediate situation, but it could be a more important result if non-binary design matrices are involved.

3.2 Strong overlap condition and rescaled likelihood

Existence and uniqueness conditions to solve (3.1) are well-known [22], they consist in an overlap condition a little bit stronger than the one given by Silvapulle. In fact, we need the point \bar{N}^1 to be surrounded by the rows of I_0 (hereafter denoted by the letter I). We give this result in the framework of the binary problem (simpler than Owen's general case) and establish a new proof leading to a minimum relative entropy problem. From there and using duality, we build the corresponding rescaled likelihood also justified by a rescaling on the initial likelihood.

Theorem 3.2. *There exists a unique finite solution to the limit imbalanced BLR problem if and only if there exists $\lambda \in (\mathbb{R}_+^*)^q$ such that $I^T \lambda = \bar{N}^1$ and $\sum_{i=1}^q \lambda_i = 1$. (If present, the null row (such that $I_i = (0, \dots, 0)$) is removed².)*

Remark 3.1. *If $\bar{N}^1 \in (\mathbb{Q}^*)^p$, the condition $I^T \lambda = \bar{N}^1$ is equivalent to the existence of $n_0^* \in (\mathbb{N}^*)^{q_0}$ such that $I_0^T n_0^* = N^1$ (see the proof of Theorem 2.1). In other words, N^1 should be attained by a certain observation of all the rows in I_0 . This equivalence is no more true if \bar{N}^1 has irrational components.*

Our proof of this theorem is based on the following three lemmas.

Lemma 3.1. *The log-sum-exp function $h : \mathbb{R}^q \mapsto \mathbb{R}$, defined by $h(z) = \log(\sum_{i=1}^q e^{z_i})$ is a convex, continuous, increasing function on \mathbb{R}^q . The function $f : \mathbb{R}^p \mapsto \mathbb{R}$, $f(\beta) = \log(\sum_{i=1}^q \bar{n}_i^0 e^{(I\tilde{\beta})_i})$ is continuous and convex on \mathbb{R}^p .*

Proof. Function h has a positive semi-definite Hessian and is then convex. Furthermore for all $y, z \in \mathbb{R}^q$ such that $y_i \leq z_i$, $i \in \{1, \dots, q\}$, we have $h(y) \leq h(z)$ and the function is increasing on \mathbb{R}^q . The composition with an affine mapping preserves continuity and convexity. Thus, with $z = I\tilde{\beta} + b$ and $\bar{n}^0 = e^b$ we obtain a convex continuous $f(\tilde{\beta}) = h(I\tilde{\beta} + b)$ and $\text{dom } f = \mathbb{R}^p$. \square

Lemma 3.2. *The function $f : \mathbb{R}^p \mapsto \mathbb{R}$, $f(\tilde{\beta}) = \log(\sum_{i=1}^q \bar{n}_i^0 e^{(I\tilde{\beta})_i})$ is strictly convex on \mathbb{R}^p .*

Proof. The Hessian H of $h : \mathbb{R}^q \mapsto \mathbb{R}$, $h(z) = \log(\sum_{i=1}^q e^{z_i})$, is the following:

$$H_{ij} = \delta_{ij} \frac{e^{z_i}}{\sum_{k=1}^q e^{z_k}} - \frac{e^{z_i}}{\sum_{k=1}^q e^{z_k}} \frac{e^{z_j}}{\sum_{k=1}^q e^{z_k}}, \quad i, j \in \{1, \dots, q\}.$$

For all $v = (v_1, \dots, v_q)^T \in \mathbb{R}^q$, we have

$$\sum_{i,j=1}^q v_i H_{ij} v_j = \frac{(\sum_{k=1}^q e^{z_k} v_k^2)(\sum_{k=1}^q e^{z_k}) - (\sum_{k=1}^q e^{z_k} v_k)^2}{(\sum_{k=1}^q e^{z_k})^2},$$

which is non-negative due to the Cauchy-Schwarz inequality. This expression is equal to zero if and only if there exists $\lambda \in \mathbb{R}$ such that $e^{z_k} v_k^2 = \lambda e^{z_k}$, $\forall k \in \{1, \dots, q\}$. Thus, only in the constant direction $z_k(t) = t + z_k(0)$, $k \in \{1, \dots, q\}$, $t \in \mathbb{R}$, the function h is affine, in any others, this function is strictly convex.

Suppose that there exists a family of parameters $F_a = \{\tilde{\beta}(t) \in \mathbb{R}^p, t \in [0, a], a > 0\}$ such that $z(t) = I\tilde{\beta}(t) + b = t + z(0)$ and $e^b = \bar{n}^0$. This means that along the path described by $\tilde{\beta}(t)$ the function f is affine. We obtain $I(\tilde{\beta}(t) - \tilde{\beta}(0)) = t$ and with $t \neq 0$, we have $\gamma = (-t, \tilde{\beta}(t) - \tilde{\beta}(0))^T \in \mathbb{R}^{p+1} \setminus \{0\}^{p+1}$ such that $I\gamma = 0$. This is impossible because the matrix I is of full rank, which proves the lemma. \square

²in order to have non-zero coefficients λ as for the overlap condition in Theorem 2.1.

We present a corollary to a theorem on the Legendre-Fenchel transform of convex composite functions exposed in [14].

Lemma 3.3. *If functions $g_i : \mathbb{R}^p \mapsto \mathbb{R}$, $i \in \{1, \dots, q\}$ are convex and continuous with $\text{dom } g_i = \mathbb{R}^p$ and $h : \mathbb{R}^q \mapsto \mathbb{R}$ is convex, continuous and increasing with $\text{dom } h = \mathbb{R}^q$, then the convex conjugate of $h(g_1, \dots, g_q)$ is given by*

$$[h(g_1, \dots, g_q)]^*(m) = \min_{\substack{\alpha_1 \geq 0, \dots, \alpha_q \geq 0 \\ m_1 + \dots + m_q = m}} \left(h^*(\alpha_1, \dots, \alpha_q) + \sum_{i=1}^q \alpha_i g_i^*\left(\frac{m_i}{\alpha_i}\right) \right),$$

with $m \in (\mathbb{R}^p)^T$.

Proof of the theorem. Let us define the function F_m such that

$$F_m : \begin{cases} \mathbb{R}^p & \rightarrow \mathbb{R}, \\ \tilde{\beta} & \mapsto m \cdot \tilde{\beta} - \log(\sum_i \bar{n}_i^0 e^{(\mathbf{I}\tilde{\beta})_i}). \end{cases}$$

F_m is differentiable on \mathbb{R}^p and the first order equations

$$\frac{\partial}{\partial \tilde{\beta}_j} \left(\sum_{i=1}^p m_i \cdot \tilde{\beta}_i - f(\tilde{\beta}) \right) = 0, \quad j \in \{1, \dots, p\},$$

are equal to the system (3.1) with $\bar{N}^1 = m^T$. Suppose that there exist two solutions $\gamma_1, \gamma_2 \in \mathbb{R}^p$ to (3.1), then $\nabla F_{\bar{N}^1}(\gamma_1) = \nabla F_{\bar{N}^1}(\gamma_2) = 0$. However, the function F_m is strictly concave as the sum of a concave function and a strictly concave function (see Lemma 3.2). Consequently, the solution γ to $\nabla F_{\bar{N}^1}(\gamma) = 0$ is unique and $\gamma_1 = \gamma_2$.

We now introduce the convex conjugate of the function f :

$$f^* : \begin{cases} (\mathbb{R}^p)^T & \mapsto \mathbb{R}, \\ m & \mapsto \sup_{\tilde{\beta} \in \mathbb{R}^p} \left(m \cdot \tilde{\beta} - \log(\sum_i \bar{n}_i^0 e^{(\mathbf{I}\tilde{\beta})_i}) \right) = \sup_{\tilde{\beta} \in \mathbb{R}^p} \left(F_m(\tilde{\beta}) \right). \end{cases}$$

We will prove that the three following sets are identical

$$A = \left\{ m \in (\mathbb{R}^p)^T \mid \exists \tilde{\beta} \in \mathbb{R}^p, \mathbf{I}^T \left(\frac{\bar{n}^0 e^{\mathbf{I}\tilde{\beta}}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\tilde{\beta})_i}} \right) = m^T \right\},$$

$$B = \left\{ m \in (\mathbb{R}^p)^T \mid f^*(m) < +\infty \right\},$$

$$C = \left\{ m \in (\mathbb{R}^p)^T \mid \exists \lambda \in (\mathbb{R}_+^*)^q, m = \lambda^T \mathbf{I}, \sum_{i=1}^q \lambda_i = 1 \right\}.$$

i) $A \subset B$. If $m_0 \in A$ there exists $\tilde{\beta} \in \mathbb{R}^p$ solution to (3.1), moreover, $\nabla F_{m_0}(\tilde{\beta}) = 0$ and $f^*(m_0) = F_{m_0}(\tilde{\beta})$ because of the strict concavity of F_{m_0} . Thus $m_0 \in B$.

ii) $B \subset C$. We use the Lemma 3.3 with $g_i(\tilde{\beta}) = (\mathbf{I}\tilde{\beta})_i + b_i$ and h the log-sum-exp function verifying the necessary conditions (Lemma 3.1). We have the convex conjugate $g_i^*(u_i) = -b_i$ if $u_i = \mathbf{I}_i$ and $+\infty$ elsewhere (we do not consider the presence of a null row $\mathbf{I}_i = (0, \dots, 0)$). The only way to obtain a finite result is to impose the constraint $u_i = \frac{m_i}{\alpha_i} = \mathbf{I}_i$ for all $i \in \{1, \dots, q\}$. Therefore, knowing that

$$h^*(\alpha_1, \dots, \alpha_q) = \begin{cases} \sum_{i=1}^q \alpha_i \log(\alpha_i) & \text{if } \alpha_1 \geq 0, \dots, \alpha_q \geq 0, \alpha_1 + \dots + \alpha_q = 1, \\ +\infty & \text{otherwise,} \end{cases}$$

we have

$$\begin{aligned}
f^*(m) &= [h(g_1, \dots, g_q)]^*(m) = \min_{\substack{\alpha_1 \geq 0, \dots, \alpha_q \geq 0 \\ \alpha_1 + \dots + \alpha_q = 1 \\ \alpha_1 \mathbf{I}_1 + \dots + \alpha_q \mathbf{I}_q = m}} \left(\sum_{i=1}^q \alpha_i \log(\alpha_i) + \sum_{i=1}^q \alpha_i (-b_i) \right) \\
&= \min_{\substack{\alpha_1 \geq 0, \dots, \alpha_q \geq 0 \\ \alpha_1 + \dots + \alpha_q = 1 \\ \alpha_1 \mathbf{I}_1 + \dots + \alpha_q \mathbf{I}_q = m}} \left(\sum_{i=1}^q \alpha_i \log\left(\frac{\alpha_i}{\bar{n}_i^0}\right) \right).
\end{aligned}$$

We minimize a Kullback–Leibler divergence between two distributions under linear constraints. If one of the α_i is zero, $0g_i^*\left(\frac{m_i}{0}\right) = \sigma_{\text{dom } g_i}(m_i) = 0$ if $m_i = 0$ elsewhere $+\infty$ (see [14]) and the previous equalities remain true with $m_i = \alpha_i \mathbf{I}_i$. The KKT conditions of this problem impose the constraint $\alpha_i > 0$ for all $i \in \{1, \dots, q\}$. Thus,

$$f^*(m) = [h(g_1, \dots, g_q)]^*(m) = \min_{\substack{\alpha_1 > 0, \dots, \alpha_q > 0 \\ \alpha_1 + \dots + \alpha_q = 1 \\ \alpha_1 \mathbf{I}_1 + \dots + \alpha_q \mathbf{I}_q = m}} \left(\sum_{i=1}^q \alpha_i \log\left(\frac{\alpha_i}{\bar{n}_i^0}\right) \right). \quad (3.4)$$

This minimum exists: this is a linear restriction to a convex and continuous function in a simplex and therefore $B = \text{dom } f^* \subset C$.

iii) $C \subset B \subset A$. If $m_0 \in C$, then there exists $\lambda \in (\mathbb{R}_+^*)^q$ such that $\sum_{i=1}^q \lambda_i = 1$ and $\lambda^T \mathbf{I} = m_0$ so that

$$f^*(m_0) = \sup_{\tilde{\beta} \in \mathbb{R}^p} \left(m_0 \cdot \tilde{\beta} - \log\left(\sum_i \bar{n}_i^0 e^{(\mathbf{I}\tilde{\beta})_i}\right) \right) = \sup_{\tilde{\beta} \in \mathbb{R}^p} \left(\log\left(\frac{e^{\sum_i \lambda_i (\mathbf{I}\tilde{\beta})_i}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\tilde{\beta})_i}}\right) \right).$$

If the supremum is reached, there is a maximizing element $\gamma \in \mathbb{R}^p$ and $m_0 \in B$, this element is the solution to the system (3.1) and thus $m_0 \in A$. To state this result, it is enough to have $-F_{m_0}$ coercive. Let $\epsilon \in \mathbb{R}^p \setminus \{0\}^p$ be an arbitrary vector and $\tilde{\beta} = x\epsilon$ with $x \in \mathbb{R}$. Then,

$$F_{m_0}(x\epsilon) = \log\left(\frac{e^{\sum_i \lambda_i (\mathbf{I}\epsilon)_i x}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\epsilon)_i x}}\right) = \log\left(\frac{e^{\sum_i \lambda_i \omega_i x}}{\sum_i \bar{n}_i^0 e^{\omega_i x}}\right),$$

with $\omega = \mathbf{I}\epsilon$. Notice that the vector ω can not satisfy the relations $\omega_1 = \dots = \omega_p$ because \mathbf{I} is of full rank. Thus, if $W = \max_{i \in \{1, \dots, q\}} (\omega_i)$, we have

$$F_{m_0}(x\epsilon) = \log\left(\frac{e^{(\sum_i (\lambda_i \omega_i) - W)x}}{\sum_i \bar{n}_i^0 e^{(\omega_i - W)x}}\right),$$

and $\sum_i (\lambda_i \omega_i) - W < 0$ because $\sum \lambda_i = 1$. Therefore, with $\Omega = \{i \in \{1, \dots, p\} \mid \omega_i = W\}$

$$\lim_{x \rightarrow +\infty} e^{(\sum_i (\lambda_i \omega_i) - W)x} = 0, \quad \lim_{x \rightarrow +\infty} \sum_i \bar{n}_i^0 e^{(\omega_i - W)x} = \sum_{i \in \Omega} \bar{n}_i^0 > 0,$$

which proves that the function $-F_{m_0}$ is coercive when $m_0 \in C$ and achieves the proof. \square

The expression $f^*(m)$ in (3.4) is the minimization of a relative entropy between the class 0 distribution and a kind of ghost class 1 distribution (built on the $\mathbf{I} = \mathbf{I}_0$ design matrix). With the duality property, we can introduce a new likelihood. The following proposition leads to the same "limit" likelihood and justifies the use of the adjective "rescaled". Indeed:

Proposition 3.2. *The limit imbalanced equations arise from the following rescaled likelihood:*

$$L^*(\tilde{\beta} | \mathbf{I}_0, \mathbf{I}_1, \bar{n}^0, n^1) = \prod_{j=1}^{q_1} \left(\frac{e^{(\mathbf{I}_1 \tilde{\beta})_j}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \right)^{n_j^1}.$$

Proof. With the initial likelihood

$$L(\beta | \mathbf{I}_0, \mathbf{I}_1, n^0, n^1) = \prod_{j=1}^{q_1} \left(\frac{e^{(\mathbf{I}_1 \beta)_j}}{1 + e^{(\mathbf{I}_1 \beta)_j}} \right)^{n_j^1} \prod_{j=1}^{q_0} \left(1 + e^{(\mathbf{I}_0 \beta)_j} \right)^{-n_j^0},$$

and the relation $\exp(\beta_0) = \frac{|n^1|}{\sum_{i=1}^{q_0} n_i^0 e^{(\mathbf{I}_0 \beta)_i} + \sum_{i=1}^{q_1} n_i^1 e^{(\mathbf{I}_1 \beta)_i}} = \frac{|n^1|}{|n^0|} C\left(\frac{|n^1|}{|n^0|}\right)$ (see 3.3), we obtain the following expression for the likelihood, using notation $x = \frac{|n^1|}{|n^0|}$:

$$x^{|n^1|} \prod_{j=1}^{q_1} \left(e^{(\mathbf{I}_1 \tilde{\beta})_j} C(x) \right)^{n_j^1} \prod_{j=1}^{q_1} \left(\frac{1}{1 + x C(x) e^{(\mathbf{I}_1 \tilde{\beta})_j}} \right)^{n_j^1} \prod_{j=1}^{q_0} \left(\frac{1}{1 + x C(x) e^{(\mathbf{I}_0 \tilde{\beta})_j}} \right)^{n_j^0}.$$

We consider that $|n^0|$ is large enough to consider the limit (with $|n^1|$ fixed) " $x \rightarrow 0$ " and to make the approximations

$$\left(e^{(\mathbf{I}_1 \tilde{\beta})_j} C(x) \right)^{n_j^1} \rightarrow \left(\frac{e^{(\mathbf{I}_1 \tilde{\beta})_j}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \right)^{n_j^1}, \quad \prod_{j=1}^{q_1} \left(\frac{1}{1 + x C(x) e^{(\mathbf{I}_1 \tilde{\beta})_j}} \right)^{n_j^1} \rightarrow 1,$$

and

$$\prod_{j=1}^{q_0} \left(\frac{1}{1 + x C(x) e^{(\mathbf{I}_0 \tilde{\beta})_j}} \right)^{n_j^0} \rightarrow e^{-|n^1|},$$

thus,

$$L(\beta) \times \left(\frac{|n^0|}{|n^1|} \right)^{|n^1|} e^{|n^1|} \rightarrow L^*(\tilde{\beta}),$$

as β_0 tends to minus infinity because $\frac{|n^0|}{|n^1|} \rightarrow +\infty$. \square

The reader can see an analogy in physics with the existence of different scales of modelization. For example, the discrete microscopic N-body problem changed into the mesoscopic Boltzmann equation using the Boltzmann-Grad limit. See the book [25] for further information on hydrodynamic limits.

This new likelihood makes now possible to consider a wide range of problems, related to variance reduction using simple prior penalties (Subsection 3.4) or regularization (Section 4).

3.3 The relative entropy dual problem

With a likelihood and an entropy, we benefit from two points of view in order to numerically estimate the regression coefficients. The classical approach using a Newton-Raphson algorithm associated to the likelihood can be challenged by other algorithms on the primal or dual problems as described in [20] and [30] for classical logistic regression. We present here the dual problem and its link with initial regression coefficients. We leave the numerical analysis to another study.

Proposition 3.3. *The regression coefficients of the limit imbalanced regression are given by the formulae*

$$\hat{\beta} = (P^T P)^{-1} P^T \log \left(\frac{n^*}{\bar{n}^0} e^{-A} \right),$$

where n^* is the probability distribution solving a relative entropy problem with linear constraints

$$n^* = \underset{\substack{\alpha_1 > 0, \dots, \alpha_q > 0 \\ \alpha_1 + \dots + \alpha_q = 1 \\ \mathbf{I}^T \alpha = \bar{N}^1}}{\operatorname{argmin}} DL(\alpha || \bar{n}^0),$$

$$A = \sum_{i=1}^q n_i^* \log \left(\frac{n_i^*}{\bar{n}_i^0} \right) \text{ and } P = \mathbf{I} - M \text{ with } P_{ij} = \mathbf{I}_{ij} - \bar{N}_j^1.$$

Proof. With the existence of a unique solution (see Subsection 3.2), there exists a solution $n^* \in (\mathbb{R}_+^*)^q$ such that $\mathbf{I}^T n^* = \bar{N}^1$, $\sum_i n_i^* = 1$, and

$$\bar{N}^1 \cdot \tilde{\beta} - \log \left(\sum_i \bar{n}_i^0 e^{(\mathbf{I} \tilde{\beta})_i} \right) = \sum_{i=1}^q n_i^* \log \left(\frac{n_i^*}{\bar{n}_i^0} \right) = A.$$

Then, using equations (3.1) we obtain

$$\mathbf{I}^T n^* = \mathbf{I}^T (\bar{n}^0 e^A e^{P \tilde{\beta}}).$$

Let H be an orthogonal matrix to I (the previous relation remains true with I instead of \mathbf{I}) and $\gamma \in \mathbb{R}^{q-p-1}$, such that we can remove \mathbf{I} to obtain the relation

$$n^* + H\gamma = \bar{n}^0 e^A e^{P \tilde{\beta}},$$

hence,

$$-\sum_{i=1}^q n_i^* \log \left(\frac{n_i^*}{\bar{n}_i^0} \right) + \log \left(\frac{n_k^* + (H\gamma)_k}{\bar{n}_k^0} \right) = \sum_{j=1}^p (\mathbf{I}_{kj} - \bar{N}_j^1) \tilde{\beta}_j, \quad k \in \{1, \dots, q\}.$$

Summing all these relations with weights $n_k^* + (H\gamma)_k$, using the fact that $\sum_{k=1}^q (H\gamma)_k = 0$, gives

$$\sum_{i=1}^q n_i^* \log \left(\frac{n_i^*}{\bar{n}_i^0} \right) - \sum_{k=1}^q (n_k^* + (H\gamma)_k) \log \left(\frac{n_k^* + (H\gamma)_k}{\bar{n}_k^0} \right) = 0.$$

Due to convexity of the Kullback-Leibler divergence, we have a unique minimum obtained (by definition of n^*) at $\gamma = 0$. Therefore

$$P \tilde{\beta} = \log \left(\frac{n^*}{\bar{n}^0} e^{-A} \right),$$

and the result is proved if P is of full rank. Suppose that this is not the case. Then, there exists $\lambda \in \mathbb{R}^p \setminus \{0\}^p$ such that $P\lambda = (\mathbf{I} - M)\lambda = 0$, therefore $\mathbf{I}\lambda = C$ with C a vector with identical components all equal to $\sum_{i=1}^p \bar{N}_i^1 \lambda_i$. Consequently, the matrix I is no more of full rank, which is, by definition of $I = I_0$, impossible. \square

3.4 Priors for variance reduction

The rare events structure of class imbalance goes hand in hand with the problem of precision for estimates. A classical solution consists in introducing an a priori distribution in a Bayesian context. This can be done using a Jeffreys non-informative prior [15] allowing both first order bias removal and variance shrinkage [7]. Thus, we have to maximize the expression

$$L_J^*(\tilde{\beta} | \mathbf{I}_0, \mathbf{I}_1, \bar{n}^0, n^1) = \prod_{j=1}^{q_1} \left(\frac{e^{(\mathbf{I}_1 \tilde{\beta})_j}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}_0 \tilde{\beta})_i}} \right)^{n_j^1} \times |\mathcal{I}(\tilde{\beta})|^{1/2},$$

with $|\mathcal{I}|$ the determinant of the Fisher information matrix. This approach is implemented in the R package *logistf* for logistic regressions. In the imbalanced case, we search for a method conserving the shape of the limit equations and achieving at the same time variance reduction: we choose the following approximation

$$\frac{1}{2} \log(|\mathcal{I}(\tilde{\beta})|) \approx \frac{1}{2} \sum_{i=1}^p \log \left(\frac{\mathbf{I}_i^T (\bar{n}^0 e^{\mathbf{I} \tilde{\beta}})}{\sum_j \bar{n}_j^0 e^{(\mathbf{I} \tilde{\beta})_j}} - \left(\frac{\mathbf{I}_i^T (\bar{n}^0 e^{\mathbf{I} \tilde{\beta}})}{\sum_j \bar{n}_j^0 e^{(\mathbf{I} \tilde{\beta})_j}} \right)^2 \right),$$

supposing an absence of correlation between predictors in a random design framework (see Section 4). With this hypothesis, we derive first order equations

$$N_i^1 - |n^1| \frac{\mathbf{I}_i^T \bar{n}^0 e^{\mathbf{I} \tilde{\beta}}}{\sum_j \bar{n}_j^0 e^{(\mathbf{I} \tilde{\beta})_j}} + \frac{1}{2} \left(1 - 2 \frac{\mathbf{I}_i^T (\bar{n}^0 e^{\mathbf{I} \tilde{\beta}})}{\sum_j \bar{n}_j^0 e^{(\mathbf{I} \tilde{\beta})_j}} \right) = 0, \quad i \in \{1, \dots, p\},$$

thus,

$$\mathbf{I}^T \left(\frac{\bar{n}^0 e^{\mathbf{I} \tilde{\beta}}}{\sum_j \bar{n}_j^0 e^{(\mathbf{I} \tilde{\beta})_j}} \right) = \frac{N^1 + \frac{1}{2}}{|n^1| + 1} = \bar{N}_J^1.$$

In table 3, we simulate data sets as previously done with the length for n^0 fixed to 10 and we compare estimated bias and variance for coefficient $\beta_4 = 0.25$ with three different methods: a classical logistic regression (*bias* and *sd.*), the imbalanced case with means \bar{N}_J^1 (*im. bias* and *im. sd.*) and the Jeffreys exact penalty (*J. bias* and *J. sd.*).

Table 3: Variance and bias analysis with prior distributions

β_0	-5	-4	-3	-2	-1	0
$ n^0 / n^1 $	143	52	19	7.1	2.7	1.0
sd.	0.3720	0.2173	0.1328	0.08835	0.06388	0.05645
im. sd.	0.3672	0.2165	0.1326	0.08837	0.06406	0.05677
J. sd.	0.3575	0.2141	0.1322	0.08814	0.06382	0.05641
bias	6.367e-3	4.247e-3	-8.032e-4	1.281e-3	-2.822e-4	3.048e-4
im. bias	3.284e-3	3.025e-3	-1.176e-3	1.159e-3	-3.229e-4	2.987e-4
J. bias	-1.631e-3	1.419e-3	-1.858e-3	8.354e-4	-5.028e-4	1.345e-4

Variance reduction is about 2 percents with the Jeffreys prior and the half as much its easily computable approximation in class imbalance. Bias was already small and gets a little smaller. The shrinkage of the variance is limited by the Cramér-Rao bound (see Fisher variance in table 1) and no miraculous reduction was conceivable.

In the next section, we consider path following methods to complete regularization and highlight its "simplicity" with binary data. The initial parameters being the maximum a

posteriori estimate (MAP), this estimation is a central problem of the limit imbalanced study. The benefit of the rescaled likelihood compared with the standard one is in the easy use of exponential a priori penalties. Indeed, with the penalty³

$$P(\tilde{\beta}) = \exp \left(\sum_{i=1}^{q_1} \epsilon_i \tilde{\beta}_i \right), \quad (3.5)$$

where $\epsilon \in \mathbb{R}^{q_1}$, we maintain the shape of the likelihood by only perturbing the predictor means vector \bar{N}^1 by $\frac{\epsilon}{|n^1|}$ (the MAP exists if and only if $\bar{N}^1 + \frac{\epsilon}{|n^1|}$ is surrounded by the rows of I , see theorem 3.2).

4 Lasso-type regularization

We consider that each observation \mathbf{I}_i ($i \in \{1, \dots, q\}$) is generated by a random binary vector $X_i = (X_{i1}, \dots, X_{ip})^T \in \{0, 1\}^p$ with $\mathbb{E}[X_{ij}] = b_j$, $j \in \{1, \dots, p\}$. With this modelization, we find many precise path estimators depending on the underlying correlation structure of the random design. We conclude this section by proposing a simple fast algorithm illustrated on pharmacovigilance data.

4.1 Limit lasso properties

The well-known lasso regularization consists in introducing a positive parameter λ defining the strength of a Laplace prior distribution [29]. We search for the maximum of the expression

$$\mathcal{L}(\beta, \lambda) = L^*(\beta) \times \exp \left(-\lambda \sum_{i=1}^p |\beta_i| \right),$$

which verifies the following simple first order conditions. Notice that we use, from now on, the notation β instead of $\tilde{\beta}$ to facilitate the reading.

Proposition 4.1. *The limit imbalanced BLR problem with lasso penalty leads to the system of equations*

$$\mathbf{I}^T \left(\frac{\bar{n}^0 e^{\mathbf{I}\beta}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}} \right) = \bar{N}^1 - t\nu(\beta), \quad (4.1)$$

with $t = \frac{\lambda}{|n^1|}$ and $\nu_j(\beta) = \text{sign}(\beta_j)$, if $\beta_j \neq 0$, $\nu_j(\beta) \in [-1, 1]$, if $\beta_j = 0$ for all $j \in \{1, \dots, p\}$ (ν is the subgradient of the l^1 norm).

Thus, the lasso has a clear interpretation as a shift operating on the observed proportions \bar{N}^1 . Thereafter, we often use the vector $p(t) \in \mathbb{R}^p$ defined as $p(t) = \bar{N}^1 - t\nu(\beta)$.

Proposition 4.2. *If the strong overlap condition in Theorem 3.2 is satisfied, then the function*

$$\hat{\beta} : \begin{cases} \mathbb{R} & \rightarrow \mathbb{R}^p, \\ t & \mapsto \underset{\beta \in \mathbb{R}^p}{\text{argmax}}(\mathcal{L}(\beta, t)), \end{cases}$$

is continuous for all $t \geq 0$ and there exists $T \in [0, 1]$, such that $\beta(t) = 0$, $\forall t \geq T$.

³P could be written as a probability distribution with a normalization term (the support of regression coefficients is finite).

Proof. With the positivity of \mathcal{L} , we have,

$$\operatorname{argmax}_{\beta \in \mathbb{R}^p}(\mathcal{L}(\beta, t)) = \operatorname{argmin}_{\beta \in \mathbb{R}^p}(-\log(\mathcal{L}(\beta, t))),$$

and for all $t \geq 0$, $-\log(\mathcal{L})$ is a strictly convex and coercive function in β if the strong overlap condition is satisfied (see proof of Theorem 3.2). Therefore, the function $\hat{\beta}$ is well defined for all $t \geq 0$. Furthermore, this function is continuous because of the continuity in (β, t) of $\log(\mathcal{L})$ and its strict concavity in β . The equations (4.1) with $t = 1$ have no solution if one of the components of $\nu(\beta)$ is equal to -1 or $+1$, therefore $\hat{\beta}_j(1) = 0$ for all $j \in \{1, \dots, p\}$. \square

Remark 4.1. Using the law of large numbers, the family of model parameters $\{\beta(t)\}_{t \geq 0}$ solves the system of equations

$$\frac{\mathbb{E}[X_{1j}e^{X_1\beta(t)}]}{\mathbb{E}[e^{X_1\beta(t)}]} = p_j(t), \quad j \in \{1, \dots, p\}, \quad (4.2)$$

with \mathbb{E} being the expectation operator. This previous system of equations takes the same form as in (4.1) because X_1 is a discrete random vector and therefore the path estimator $\{\beta(t)\}_{t \geq 0}$ is continuous. Notice that the function $\nu \circ \beta: \mathbb{R}^+ \rightarrow \mathbb{R}$ is also continuous in t .

4.2 Path estimators

Thanks to this previous remark, we are able to find precise analytic estimators of the path in the case of independent and orthogonal random designs. Notice that such solutions already exist in the framework of linear regression (see [29]). From now on, the strong overlap condition is considered to be always satisfied at $t = 0$.

Theorem 4.1. If the random vector X generating the observations \mathbf{I} has independent components, a precise path estimator $\{\hat{\beta}(t)\}_{t \geq 0}$ is given by the formulae

$$\hat{\beta}_j(t) = \hat{\beta}_j(0) + \log \left(\frac{1 - t \frac{\operatorname{sign}(\hat{\beta}_j)}{\bar{N}_j^1}}{1 + t \frac{\operatorname{sign}(\hat{\beta}_j)}{1 - \bar{N}_j^1}} \right), \quad j \in \{1, \dots, p\},$$

$$\text{if } t \in [0, t_{0j}], \quad t_{0j} = \bar{N}_j^1 \frac{|1 - e^{-\hat{\beta}_j(0)}|}{1 + \frac{\bar{N}_j^1}{1 - \bar{N}_j^1} e^{-\hat{\beta}_j(0)}} \quad \text{and} \quad \hat{\beta}_j(t) = 0, \quad \text{if } t > t_{0j}.$$

The coefficients $\hat{\beta}(0)$ are given by the classical MLE (solution of (2.2) without intercept) if we want to estimate the path obtained by an imbalanced logistic regression or by the MLE of the rescaled likelihood (Proposition 3.2), in this case,

$$\hat{\beta}_j(0) = \log \left(\frac{\bar{N}_j^1}{1 - \bar{N}_j^1} \frac{1 - \bar{N}_j^0}{\bar{N}_j^0} \right), \quad j \in \{1, \dots, p\}.$$

Proof. For all $j \in \{1, \dots, p\}$ we use the hypothesis of independence:

$$p_j(t) = \frac{\mathbb{E}[X_j e^{X\beta(t)}]}{\mathbb{E}[e^{X\beta(t)}]} = \frac{\mathbb{E}[X_j e^{X_j \beta_j(t)}] \mathbb{E}[\prod_{k \neq j} e^{X_k \beta_k(t)}]}{\mathbb{E}[e^{X_j \beta_j(t)}] \mathbb{E}[\prod_{k \neq j} e^{X_k \beta_k(t)}]} = \frac{\mathbb{E}[X_j e^{X_j \beta_j(t)}]}{\mathbb{E}[e^{X_j \beta_j(t)}]}$$

$$= \frac{e^{\beta_j(t)} P(X_j = 1)}{e^{\beta_j(t)} P(X_j = 1) + P(X_j = 0)} = \frac{e^{\beta_j(t)} b_j}{e^{\beta_j(t)} b_j + (1 - b_j)},$$

and the solution is

$$\beta_j(t) = \log \left(\frac{p_j(t)}{1 - p_j(t)} \right) - \log \left(\frac{b_j}{1 - b_j} \right), \quad t \in [0, |\bar{N}_j^1 - b_j|],$$

and $\beta_j(t) = 0$ if $t > |\bar{N}_j^1 - b_j|$. Indeed, $\dot{\beta}_j$ is negative in region $\beta_j > 0$ and positive in region $\beta_j < 0$. $\beta_j(0) = \log \left(\frac{\bar{N}_j^1}{1 - \bar{N}_j^1} \frac{1 - b_j}{b_j} \right)$ for a random design with independent predictors (see Appendix A.2 with $p = 1$). We replace all the b_j by the frequencies of observations \bar{N}_j^0 to obtain the estimator. \square

The orthogonal case, when the inner product between columns of the design matrix vanishes ($X_{1j}X_{1k} = 0$, $j \neq k$), is also tractable.

Theorem 4.2. *If the random design is orthogonal, we have $\mathbf{I} \in \mathcal{M}_{(p+1) \times p}(\mathbb{B})$ filled by zeros except at positions $(i+1, i)$, $i = 1, \dots, p$ and the derivative of the path estimator takes the form*

$$\dot{\beta}_i(t) = \log \left(\frac{p_i(t)}{1 - \sum_{s \in S_t} p_s(t)} \right), \quad i \in S_t, \quad t \geq 0,$$

with $\{S_t\}_{t \geq 0}$ a family of subsets of $\{1, \dots, p\}$ containing the indexes of non-zero coefficients of vector β at time $t \geq 0$. The algorithm that describes the positions of the change-points in S_t is described in the proof.

Proof. With the hypothesis of orthogonality, equations (4.2) are reduced to

$$\begin{cases} \frac{b_1 e^{\beta_1(t)}}{b_0 + b_1 e^{\beta_1(t)} + \dots + b_p e^{\beta_p(t)}} = \bar{N}_1^1 - t \nu_1(t), \\ \dots \\ \frac{b_p e^{\beta_p(t)}}{b_0 + b_1 e^{\beta_1(t)} + \dots + b_p e^{\beta_p(t)}} = \bar{N}_p^1 - t \nu_p(t), \end{cases} \quad (4.3)$$

and we obtain

$$e^{\beta_i(t)} = \frac{b_0}{b_i} \frac{p_i(t)}{1 - \sum_{j=1}^p p_j(t)}, \quad i = 1, \dots, p.$$

Let $\bar{S}_t = \{0, 1, \dots, p\} \setminus S_t$ and $\bar{S}_t^* = \bar{S}_t \setminus \{0\}$, then

$$\begin{cases} e^{\beta_i(t)} = \frac{b_0}{b_i} \frac{\bar{N}_i^1 - t \text{sign}(\beta_i)}{1 - \sum_{j=1}^p \bar{N}_j^1 + t \sum_{s \in S_t} \text{sign}(\beta_s) + t \sum_{s \in \bar{S}_t^*} \nu_s(t)}, & i \in S_t, \\ 1 = \frac{b_0}{b_i} \frac{\bar{N}_i^1 - t \nu_i(t)}{1 - \sum_{j=1}^p \bar{N}_j^1 + t \sum_{s \in S_t} \text{sign}(\beta_s) + t \sum_{s \in \bar{S}_t^*} \nu_s(t)}, & i \in \bar{S}_t^*. \end{cases} \quad (4.4)$$

After computation, we have explicit formulae for the continuous functions β_i and ν_i ($i \in \{1, \dots, p\}$):

$$\begin{cases} e^{\beta_i(t)} = \frac{b^S}{b_i} \frac{p_i(t)}{1 - \sum_{s \in S_t} p_s(t)}, & i \in S_t, \\ \nu_i(t) = \frac{1}{t} \frac{b^S \bar{N}_i^1 - b_i \bar{N}^S}{b^S} - \frac{b_i R^S}{b^S}, & i \in \bar{S}_t^*, \end{cases} \quad (4.5)$$

with $b^S = \sum_{s \in \bar{S}_t} b_s$, $\bar{N}^S = \sum_{s \in \bar{S}_t} \bar{N}_s^1$ and $R^S = \sum_{s \in S_t} \text{sign}(\beta_s)$. These functions are monotonous, we need the change-points to draw the path, that is the finite sequence of different models $\{S_t\}_{t \geq 0} = \{S_{t_0}, S_{t_1}, \dots, S_{t_m}\}$, $m \in \mathbb{N}^*$. For all $i \in \{0, \dots, m-1\}$,

$\{S_t\}_{t \in [t_i, t_{i+1}[}$ is a unique subset. If $t \in [t_i, t_{i+1}[$ and we know S_t we determine $S_{t_{i-1}}, S_{t_{i+1}}$, t_i and t_{i+1} by solving

$$\begin{cases} \beta_i(u_i) = 0 & \Leftrightarrow u_i = \frac{b^S \bar{N}_i^1 - b_i \bar{N}^S}{b^S \text{sign}(\beta_i) + b_i R^S}, & i \in S_t, \\ \nu_i(v_i^+) = 1 & \Leftrightarrow v_i^+ = \frac{b^S \bar{N}_i^1 - b_i \bar{N}^S}{b^S + b_i R^S}, & i \in \bar{S}_t^*, \\ \nu_i(v_i^-) = -1 & \Leftrightarrow v_i^- = \frac{b^S \bar{N}_i^1 - b_i \bar{N}^S}{-b^S + b_i R^S}, & i \in \bar{S}_t^*. \end{cases} \quad (4.6)$$

We define $W = \{w_i\} = \{u_i, v_j^+, v_j^-, i \in S_t, j \in \bar{S}_t^*\}$ and the two adjacent change-points are given by

$$t_{i+1} = \min_j \{w_j \mid w_j > t\} \quad \text{and} \quad t_i = \max_j \{w_j \mid w_j \leq t\}.$$

Therefore,

$$S_{t_{i+1}} = S_t \cup V_{i+1} \setminus U_{i+1} \quad \text{and} \quad S_{t_{i-1}} = S_t \cup V_i \setminus U_i,$$

with $U_i = \{j \in \{1, \dots, p\} \mid u_j = t_i\}$, $V_i = \{j \in \{1, \dots, p\} \mid v_j^+ = t_i \text{ or } v_j^- = t_i\}$.

The path can be built forward or backward. If we choose the path following approach (forward), S_{t_0} is found using the MLE of the rescaled likelihood (see Section 3) and $t_0 = 0$. In the other configuration (backward), we have $S_{t_m} = \emptyset$ and for $t > t_m$, $b^S = \bar{N}^S = 1$ and $R^S = 0$, so that $t_m = \max_{i \in \{1, \dots, p\}} |\bar{N}_i^1 - b_i|$. \square

Simulations with this type of design show that each path usually vanishes only one time (and does not reappear) and thus $m > p$ is a very rare configuration.

The opposite situation to orthogonality is inclusion. For example, if X_{12} is included in X_{11} meaning that for the observed data $\mathbf{I}_{i1} = 1$ if $\mathbf{I}_{i2} = 1$, we find an analytic description of the estimator given by the formulae

$$\hat{\beta}_1(t) = \log \left(\frac{p_1(t)}{p_2(t) - p_1(t)} \right), \quad \hat{\beta}_2(t) = \log \left(\frac{p_2(t) - p_1(t)}{1 - p_2(t)} \right), \quad \forall t \in \{u, \hat{\beta}_1(u) \hat{\beta}_2(u) \neq 0\}.$$

This solution is likely generalizable (with a design in stairs as presented in Appendix A.4), however, this case is meaningless in the analysis of spontaneous reports databases and then left aside.

We give examples of plots of path estimates compared with a standard (using L not L^*) lasso path for different imbalance strengths in appendix C. The results highlight the high quality of the analytic path estimators.

Remark 4.2. Another regularization method is called the elastic net penalization and uses, in addition to the lasso, a second penalized term of ridge (or Tikhonov) kind [31]:

$$\mathcal{L}(\beta|y) = L^*(\beta|y) \times \exp \left(-\lambda \left[\alpha \sum_{i=1}^p |\beta_i| + \frac{1-\alpha}{2} \sum_{i=1}^p \beta_i^2 \right] \right),$$

with $\alpha \in]0, 1]$. In the case of independence in random vector X , we have an explicit formula for t with respect to β :

$$t = \frac{\bar{N}^1}{\alpha \text{sign}(\hat{\beta}) + (1-\alpha)\hat{\beta}} \frac{1 - e^{-\hat{\beta}(0) + \hat{\beta}}}{1 + \frac{\bar{N}^1}{1-\bar{N}^1} e^{-\hat{\beta}(0) + \hat{\beta}}},$$

for $\hat{\beta}$ between 0 and $\hat{\beta}(0)$. The coefficients vanish when $t_0^{en} = \frac{1}{\alpha} t_0^{lasso}$. The proof of this result is a simple adaptation of the proof for the lasso in theorem 4.1.

4.3 Negative correlation structure

If the random design verifies the relations $\mathbb{E}[X_{1j}X_{1k}e^{X_{1\beta}}]\mathbb{E}[e^{X_{1\beta}}] \leq \mathbb{E}[X_{1j}e^{X_{1\beta}}]\mathbb{E}[X_{1k}e^{X_{1\beta}}]$, $\forall j \neq k$, $\forall t \geq 0$, this in-between situation of a β -dependent negative correlation between variables X_j ($j = 1, \dots, p$) is also tractable and particularly interesting in the sparse context of near-zero components for vector \bar{N}^1 ⁴. We find two estimators that surrounded the real path.

Theorem 4.3. *The path estimator in the β -dependent negative correlation case is surrounded by estimators, whose derivatives are given by*

$$\overbrace{\log\left(\frac{p_j(t)}{1 - \sum_{s \in S_t^+} p_s(t)}\right)} \leq \dot{\beta}_j(t) \leq \overbrace{\log\left(\frac{p_j(t)}{1 - \sum_{s \in S_t^-} p_s(t)}\right)}, \quad j \in S_t,$$

with $S_t^+ = \{j \in S_t \mid \text{sign}(\beta_j) > 0\}$ and $S_t^- = \{j \in S_t \mid \text{sign}(\beta_j) < 0\}$. With the rare occurrence of resurgence of a coefficient after vanishing, we neglect this possibility and we easily find the p vanishing points and thus the family of subsets $\{S_t\}_{t \geq 0}$.

Proof. We differentiate equations (4.1) with respect to t considering only the equations verifying the condition $\beta_j(t) \neq 0$, i.e. $j \in S_t$. We obtain at time t ,

$$\begin{aligned} & \dot{\beta}_j(t) \left(\frac{\mathbf{I}_j^T(\bar{n}^0 e^{\mathbf{I}\beta})}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}} - \left(\frac{\mathbf{I}_j^T(\bar{n}^0 e^{\mathbf{I}\beta})}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}} \right)^2 \right) + \\ & \sum_{k \neq j, k \in S_t} \dot{\beta}_k(t) \left(\frac{\sum_i \mathbf{I}_{ij} \mathbf{I}_{ik} \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}} - \frac{\mathbf{I}_j^T(\bar{n}^0 e^{\mathbf{I}\beta})}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}} \frac{\mathbf{I}_k^T(\bar{n}^0 e^{\mathbf{I}\beta})}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}} \right) = -\text{sign}(\beta_j), \end{aligned}$$

or written differently,

$$\dot{\beta}_j(t) p_j(t) (1 - p_j(t)) + \sum_{k \neq j, k \in S_t} \dot{\beta}_k(t) (R_{jk}(t) - p_j(t) p_k(t)) = \dot{p}_j(t), \quad (4.7)$$

where $R_{jk}(t) = \frac{\sum_i \mathbf{I}_{ij} \mathbf{I}_{ik} \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}}{\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}}$ is a t -dependent proportion of rows with a one on the columns j and k . With only negative correlations or independence between components of X , we define the matrix $F(t) \in \mathcal{M}_{r \times r}([0, 1])$ with $r = \#S_t$ as long as $R_{jk}(t) \leq p_j(t) p_k(t)$,

$$R_{jk}(t) - p_j(t) p_k(t) = (F_{jk}(t) - 1) p_j(t) p_k(t),$$

if observations \mathbf{I} give such a matrix F . We obtain $(\mathbb{I}_r - (D - F)P)\dot{\beta} = P^{-1}\dot{P}(1)$ with P a diagonal matrix filled with the elements $\{p_i(t), i \in S_t\}$. Matrix D is the correlation-track matrix containing ones at positions (j, k) if $F_{jk}(0) < 1$ and we have

$$\dot{\beta}(t) = \sum_{i=0}^{+\infty} ((D - F)P)^i P^{-1} \dot{P}(1) = P^{-1} \dot{P}(1) + \sum_{i=1}^{+\infty} ((D - F)P)^i P^{-1} \dot{P}(1),$$

so that, using the positivity of all the elements in matrix $D - F$:

$$\overbrace{\log(P(1))} - \overbrace{\log(1 - DP^+(1))} \leq \dot{\beta}(t) \leq \overbrace{\log(P(1))} - \overbrace{\log(1 - DP^-(1))},$$

⁴Spontaneous reports databases are an example of such a sparsity with negative correlation.

with P^+ the diagonal matrix filled with vector $p^+(t) = (\max(\text{sign}(\beta_i), 0))p_i(t)_{i \in S_t}$ and P^- with vector $p^-(t) = (\max(-\text{sign}(\beta_i), 0))p_i(t)_{i \in S_t}$. Finally,

$$\overbrace{\log \left(\frac{p_j(t)}{1 - (Dp^+(t))_j} \right)}^{\cdot} \leq \dot{\beta}_j(t) \leq \overbrace{\log \left(\frac{p_j(t)}{1 - (Dp^-(t))_j} \right)}^{\cdot}, \quad j \in S_t.$$

□

In presence of sparsity (small components in \bar{N}^1), $p_j(t) \ll 1 - (Dp^-(t))_j(t)$ and $p_j(t) \ll 1 - (Dp^+(t))_j(t)$, which makes previous upper and lower bounds good path estimators. The p change-points are determined step by step and the path $\hat{\beta}$ is between a lower path and an upper path.

4.4 The piecewise logarithmic approximate path

Path following algorithms [23] are competing methods with more used coordinate descent algorithms [9] [10] (see the R package *glmnet*) or (less used) LARS [5]. In the context of linear regression, there exist exact piecewise solution paths [24]. In this subsection we introduce an algorithm giving approximate paths of piecewise logarithms for limit imbalanced logistic regression.

Proposition 4.3. *The path following algorithm for limit imbalanced logistic regression by binary predictors is the following:*

- i) $i = 0$, $t_0 = 0$, $\beta(t_0)$ given.
- ii) while $\#S_{t_i} \neq 0$ do

$$t_{i+1} = t_i + \min A_i, \quad A_i = \{a_j = \frac{1 - e^{\beta_j(t_i)}}{C_i^{-1}(\frac{\text{sign}(\beta)}{p(t_i)})} \mid a_j > 0\}, \quad (4.8)$$

with C_i :

$$(C_i)_{jk}(t_i) = \left(\frac{\sum_i \mathbf{I}_{ij} \mathbf{I}_{ik} \bar{n}_i^0 e^{(\mathbf{I}\beta)_i(t_i)}}{\mathbf{I}_k^T(\bar{n}^0 e^{\mathbf{I}\beta(t_i)})} - p_k(t_i) \right), \quad j, k \in \{1, \dots, p\}.$$

The path is then:

$$\beta(t) - \beta(t_i) = \log \left(1 - C_i^{-1} \text{sign}(\beta) \frac{t - t_i}{p(t_i)} \right), \quad t \in [t_i, t_{i+1}],$$

iii)

$$S_{t_{i+1}} = S_{t_i} \cup A_i \setminus B_i,$$

Let k be in S_t :

$$A_i = \left\{ j \in \bar{S}_{t_i}^* \mid \frac{1}{t_{i+1}} \left(\bar{N}_j^1 - (\bar{N}_j^1 - t_{i+1} \text{sign}(\beta_k(t_{i+1}))) \frac{\mathbf{I}_j^T(\bar{n}^0 e^{\mathbf{I}\beta(t_{i+1})})}{\mathbf{I}_k^T(\bar{n}^0 e^{\mathbf{I}\beta(t_{i+1})})} \right) \notin]0, 1[\right\}$$

$$B_i = \left\{ j \in S_t \mid a_j = \min A_i \right\}$$

i becomes $i + 1$.

End do.

Proof. Equations (4.7) take the form $C(t)\dot{\beta} = \overbrace{\log(p(t))}^{\cdot}$ with $C(t)$ called correction matrix.

$$C_{jk}(t) = \left(\frac{\sum_i \mathbf{I}_{ij} \mathbf{I}_{ik} \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}}{\mathbf{I}_k^T (\bar{n}^0 e^{\mathbf{I}\beta})} - p_k(t) \right), \quad j, k \in \{1, \dots, p\}.$$

Between two annulations of regression coefficients along the path (t_i and t_{i+1}), we consider this matrix to be constant ($C(t_i) = C_i$). In this case,

$$\beta(t_{i+1}) - \beta(t_i) = C_i^{-1} [\log(p(t_{i+1})) - \log(p(t_i))].$$

Of course $t_0 = 0$, but the sequence of times t is unknown. However, we iteratively approximate all its values as follows. With

$$\begin{aligned} \beta(t_{i+1}) - \beta(t_i) &= C_i^{-1} \log \left(1 - \text{sign}(\beta) \frac{t_{i+1} - t_i}{p(t_i)} \right) \\ &\approx \log \left(1 - C_i^{-1} \text{sign}(\beta) \frac{t_{i+1} - t_i}{p(t_i)} \right), \end{aligned} \quad (4.9)$$

we obtain the piecewise logarithmic path:

$$\beta(t) - \beta(t_i) = \log \left(1 - C_i^{-1} \text{sign}(\beta) \frac{t - t_i}{p(t_i)} \right), \quad t \in [t_i, t_{i+1}[,$$

with

$$t_{i+1} = t_i + \min A_i, \quad A_i = \{a_j = \frac{1 - e^{\beta_j(t_i)}}{C_i^{-1} \left(\frac{\text{sign}(\beta)}{p(t_i)} \right)} \mid a_j > 0\}, \quad i = 0, \dots, p-1.$$

The sets A_i and B_i give the indexes of regression coefficients to add or remove to S_t . \square

Other approximations could be performed, for example using a second order term in the previous approximation (4.9) or $p(t)/(1 - p(t))$ instead of $p(t)$ in the logarithm. Simulation tests show that our choice seems to give better results. In this algorithm, there is almost never any resurgence of a coefficient after a first vanishing and the simplification $A_i = \emptyset$ can be done.

Remark 4.3. *This algorithm has two main computational advantages. Firstly, the sequence t is directly determined, whereas other algorithms use a regular discretization on a logarithmic scale, which increases the number of steps to build a precise path. Secondly, the sum $\sum_i \bar{n}_i^0 e^{(\mathbf{I}\beta)_i}$ does not appear in the C_i matrices, which can highly reduce the computations especially if the matrix \mathbf{I} is sparse (0.03% of ones in the French data base). Furthermore, in comparison with the cyclic coordinate descent algorithm, there is no loop at a fixed parameter λ .*

To explore the efficiency of the algorithm, we simulate data sets with different correlation structures. Model selection is often provided with the BIC [26], which requires to know the different models arising along the path. That is why we evaluate the algorithm using a simple indicator: a comparison of the sequence of coefficients in the order of emergence along the path. The indicator is p'/p if a simulation with our algorithm gives p' coefficients at the same index as in the sequence obtained by a classical lasso algorithm (coordinate descent in R package *glmnet*). The correlation coefficient (from $r = 0$ to $r = 0.9$) means that we chose initial $R_{jk} = (1 - r)b_j b_k + r \min(b_j, b_k)$.

Table 4: Path model sequence analysis										
nb/r	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
3 (i)	0.977	0.909	0.862	0.809	0.740	0.706	0.678	0.642	0.620	0.564
3 (a)	0.777	0.710	0.736	0.730	0.735	0.720	0.727	0.726	0.725	0.755
5 (i)	0.970	0.859	0.771	0.671	0.613	0.552	0.513	0.488	0.410	0.361
5 (a)	0.785	0.752	0.725	0.715	0.700	0.760	0.708	0.693	0.699	0.705
8 (i)	0.969	0.868	0.750	0.609	0.550	0.486	0.433	0.363	0.341	0.309
8 (a)	0.773	0.742	0.717	0.756	0.722	0.698	0.714	0.675	0.684	0.622

We simulate 10^3 paths for each number nb and r , nb being the number of predictors in correlation. For each path, $\beta_0 = -5$ and the 10 regression coefficients ($\beta_1, \dots, \beta_{10}$) are always the same and chosen on a regular scale between -0.5 and 0.5 .

With $nb = 3, 5$ or 8 correlated predictors over the 10 used, the exact solution with assumption of independence (i) deteriorates with the increase in correlation (r), which is (almost) not the case if we use the approwimative algorithm (a). Notice that, with a result around 0.8, the approximate path is often very close to the exact one, this is due to close t_i terms in the t sequence (see (4.8)).

4.5 Path reconstruction with the French spontaneous reports database

We illustrate the efficiency of the limit path construction by piecewise logarithmic functions on the French spontaneous reports database. We look at two examples, a first one with no evidence of correlation and a second with strong correlations. The database contains about 330000 reports in 2016 and the imbalance is high or very high for all the adverse effects [2]. In the following graphs, the dotted lines represent results obtained by our algorithm, the solid ones result from the classical *glmnet* package.




Figure 1: Example 1. Path of regression coefficients with respect to $\log(t)$. The correlation between predictors is weak and the path given by the exact solution with independent predictors is a good approximation of the result.

The Figure 2 shows common features encountered with other examples. The path of the exponential of the coefficients shapes a set of piecewise linear functions and the




Figure 2: Example 2. We plotted the exponential of the beta coefficients with respect to t . On the left side by the assumption of no correlation, on the right side with the correction due to the approximate path algorithm. In presence of strong correlations, the algorithm of the approximate path (described in previous subsection) gives better adjusted graphs.

algorithm remains efficient even if the number of predictors is high (150 for examples). It seems that there is no case of a path with a curve reappearing after a first canceling (due to a strong correlation between predictors with opposite signs of initial coefficients). Thus, the sets B_i in the algorithm do not have to be determined.

We notice that the accuracy of this path following algorithm can easily be increased by adding intermediate steps (in variable t). The main computing limitation being the matrix inversion, one could study the inner product (Gram) matrix for class 0 and reorder rows and columns to reveal patterns and form a block diagonal matrix. These blocks could result from a statistical study of the Gram matrix⁵ (finding the pairwise independent predictors) as well as from pharmacological assumptions (medical treatments also shape patterns). Thus, computational costs become a marginal problem and one can concentrate on the bias correction by adding priors related to temporal bias, under-reporting or the introduction of similarity modifying the R matrix⁶.

5 Conclusion and perspectives

The central novelty of this work is the introduction of a rescaled likelihood for the limit imbalanced logistic regression problem. The expression of this likelihood could have some connexions with the well-known likelihoods of the self-controlled case series method [28] and of the proportional hazards model [4] used in epidemiology.

Most results exposed for binary data can be extended to other data types. However, simulations have been done only with binary data, having in mind the underlying applied problem of pharmacovigilance. The new estimate is always very close to the initial MLE because data are located on the vertexes of the hypercube and then one another "close". A convergence study of all possible existing algorithms for the primal and dual problems could be performed with different class imbalances and an evaluation of the first order term.

⁵To that end, see the literature of the block clustering problem [11].

⁶With a similarity matrix $S \in \mathcal{M}_{p \times p}([0, 1])$, the similarity is defined as follows: coefficient R_{jk} becomes $(1 - S_{jk})R_{jk} + S_{jk} \min(p_j, p_k)$.

The variance reduction is a central issue that has to be treated in a Bayesian framework. Whereas the prior to add in the standard logistic regression is unclear, the rescaled likelihood takes a well-adapted form for exponential priors. We considered model selection using the BIC and the lasso to answer this question. Due to binary data, the lasso regularization problem became easier to understand in our limit imbalanced case: we found many precise estimators. Piecewise logarithmic approximate paths are built by an effective path following procedure which determines step by step the vanishing time of each path, do not use any loops as in coordinate descent algorithms and computes expressions only involving non-zero data. Moreover, this algorithm can take into account the correlation structure between predictors to further shrink computational costs. The values for \bar{N}^1 , \bar{n}^0 and for matrix R could be shifted in order to incorporate absolute bias, temporal bias, under-reporting and similarity or correlation corrections.

6 A pharmacovigilance project?

Within this paper, we have had in mind the pharmacovigilance context as this work was carried out in parallel of a one-year engineering job at the French National Institute of Health and Medical Research. We hope this article could contribute a little to the development of mathematical tools for pharmacovigilance purposes. The science of drug safety at a postmarketing level is nearly non-existent in France as in many other countries: the reporting process of spontaneous reports is inadequate and resulting databases are badly processed with unadapted tools. Public health scandals related to medication are steadily increasing and the spotlights are turned towards big pharmaceutical companies while patient associations should firstly require public authorities to establish a modern drug safety structure. To that end, the statistical community has a major role to play by proposing trustworthy decision-support tools, opposing science to political and financial influences. Creating a useful tool was the guideline of this present work and the author hopes that other mathematicians will embrace the direction initiated by this article.

We would like to conclude by giving our opinion about the work that remains to be done to obtain an operational tool (in five points), hoping that it will inspire epidemiologists.

1) Building priors related to bias (temporal bias, under-reporting...) with the help of pharmacologists. 2) Developing new path following algorithms of different complexity and accuracy levels using (if possible) block diagonalization. 3) Introducing simple indicators to control the quality of the limit approximation. 4) Working on path visualization and new indicators (that are not thresholds). 5) Evaluating the obtained tool in the hands of pharmacologists (the use of reference sets is, to our mind, inadequate).

Acknowledgment

I would like to deeply thank Laetitia Comminges from the Paris-Dauphine University for relevant comments that greatly improved the manuscript.

A Exact solutions

We give a collection of examples consisting in simple solutions of the equation (2.3).

A.1 No intercept

If there is no intercept and no interaction between the regressors, the matrix \mathbf{I} equals the identity matrix \mathbb{I}_p and

$$\beta_1 = \log\left(\frac{n_1^1}{n_1^0}\right), \quad \dots, \quad \beta_p = \log\left(\frac{n_p^1}{n_p^0}\right).$$

If one row contains other ones, the inverse matrix is the same matrix with the added ones transformed into its opposite.

A.2 Intercept

If the square matrix \mathbf{I}_{p+1} is the following

$$\mathbf{I}_{p+1} = \begin{pmatrix} 1 & 0 & \dots & \dots & 0 \\ \vdots & 1 & 0 & \dots & 0 \\ \vdots & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 1 & 0 & \dots & 0 & 1 \end{pmatrix}, \quad \text{then} \quad \mathbf{I}_{p+1}^{-1} = \begin{pmatrix} 1 & 0 & \dots & \dots & 0 \\ -1 & 1 & 0 & \dots & 0 \\ \vdots & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ -1 & 0 & \dots & 0 & 1 \end{pmatrix}$$

for the inverse matrix, so that the β coefficients take the form

$$\beta_0 = \log\left(\frac{n_0^1}{n_0^0}\right), \quad \beta_1 = \log\left(\frac{n_1^1 n_0^0}{n_1^0 n_0^1}\right), \quad \dots, \quad \beta_p = \log\left(\frac{n_p^1 n_0^0}{n_p^0 n_0^1}\right).$$

A.3 Intercept with one correlation

The first row of the following matrix

$$\mathbf{I}_{p+1} = \begin{pmatrix} 1 & ? & \dots & \dots & ? \\ \vdots & 1 & 0 & \dots & 0 \\ \vdots & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 1 & 0 & \dots & 0 & 1 \end{pmatrix}$$

defines the set $K = \{j \in \{0, \dots, p\}, \mathbf{I}_{1j} = 1\}$. The case $\#J = 2$ is left out because it does not coincide with a non-singular matrix. The case $\#J = 1$ corresponds to the previous example. The easiest way to solve this example is to look at initial equations (2.1). We write down the $p + 1$ equations, where only the first one has a different form:

$$\sum_{i=0}^p (n_i^1 - n_i^0) = \sum_{i=1}^p (n_i^1 + n_i^0) \tanh\left(\frac{\beta_0 + \beta_i}{2}\right) + (n_0^1 + n_0^0) \tanh\left(\frac{1}{2} \sum_{j \in K} \beta_j\right)$$

and for $k \in \{1, \dots, p\}$,

$$(n_k^1 - n_k^0) + 1_{k \in K} (n_0^1 - n_0^0) = (n_k^1 + n_k^0) \tanh\left(\frac{\beta_0 + \beta_k}{2}\right) + 1_{k \in K} (n_0^1 + n_0^0) \tanh\left(\frac{1}{2} \sum_{j \in K} \beta_j\right). \quad (\text{A.1})$$

Subtracting all the p equations to the first one, we obtain

$$(1 - \#K)(n_0^1 - n_0^0) = (1 - \#K)(n_0^1 + n_0^0) \tanh\left(\frac{1}{2} \sum_{j \in K} \beta_j\right),$$

which can be used to simplify equations (A.1) into

$$(n_k^1 - n_k^0) = (n_k^1 + n_k^0) \tanh\left(\frac{\beta_0 + \beta_k}{2}\right).$$

Finally, we have

$$\exp\left(\sum_{j \in K} \beta_j\right) = \frac{n_0^1}{n_0^0} \quad \text{and} \quad \exp(\beta_0 + \beta_k) = \frac{n_k^1}{n_k^0}, \quad k \in \{1, \dots, p\},$$

so that we deduce the following closed form for the β coefficients

$$\begin{aligned} \exp(\beta_0) &= \left(\frac{n_0^0}{n_0^1} \prod_{j \in J, j \neq 0} \left(\frac{n_j^1}{n_j^0} \right) \right)^{\frac{1}{\#J-2}}, \\ \exp(\beta_i) &= \frac{n_i^1}{n_i^0} \left(\frac{n_0^1}{n_0^0} \prod_{j \in J, j \neq 0} \frac{n_j^0}{n_j^1} \right)^{\frac{1}{\#J-2}}, \quad i \in \{1, \dots, p\}. \end{aligned}$$

Notice that the regression coefficients behave in a very unpredictable way. It is sufficient to see that on an example with $p = 2$ and $\#J = 3$. The matrix \mathbf{I} is

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix},$$

and we have

$$\exp(\beta_0) = \frac{n_0^0 n_1^1 n_2^1}{n_0^1 n_1^0 n_2^0}, \quad \exp(\beta_1) = \frac{n_0^1 n_2^0}{n_0^0 n_2^1}, \quad \exp(\beta_2) = \frac{n_0^1 n_1^0}{n_0^0 n_1^1}.$$

The first intuition is to think that coefficients β_1 and β_2 do depend on the couples (n_1^0, n_1^1) and (n_2^0, n_2^1) respectively, but it is not the case!

A.4 Stairs

With

$$\mathbf{I}_{p+1} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 1 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 1 & \dots & 1 & 1 \end{pmatrix}, \quad \text{we find} \quad \mathbf{I}_{p+1}^{-1} = \begin{pmatrix} 1 & 0 & \dots & \dots & 0 \\ -1 & \ddots & \ddots & & \vdots \\ 0 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{pmatrix}$$

and then

$$\beta_0 = \log\left(\frac{n_0^1}{n_0^0}\right), \quad \beta_i = \log\left(\frac{n_i^1 n_{i-1}^0}{n_i^0 n_{i-1}^1}\right), \quad i \in \{1, \dots, p\}.$$

B Proof of proposition 3.1.

The result is proved with a succession of Taylor expansions of degree 1 or 2 in $1/s$. We use

$$\tanh\left(\frac{x}{2}\right) = -1 + 2e^x - 2e^{2x} + o(e^{2x}),$$

then with (2.2), we have

$$I_1^T n^1 = I_1^T n^1 e^{I_1 \beta} + I_0^T n^0 e^{I_0 \beta} - \left(I_1^T n^1 e^{2I_1 \beta} + I_0^T n^0 e^{2I_0 \beta} \right) + o(|n| e^{2\beta_0}). \quad (\text{B.1})$$

The first equation of this system gives

$$|n^1| = \sum n_i^1 e^{(I_1 \beta)_i} + \sum n_i^0 e^{(I_0 \beta)_i} - \left(\sum n_i^1 e^{2(I_1 \beta)_i} + \sum n_i^0 e^{2(I_0 \beta)_i} \right) + o(|n| e^{2\beta_0}),$$

therefore, using notations introduced in the proposition,

$$e^{2\beta_0} (n_2^1 + s n_2^0) - e^{\beta_0} (n_1^1 + s n_1^0) + 1 - o\left(\frac{1}{s}\right) = 0,$$

and

$$e^{\beta_0} = \frac{1}{2} \frac{n_1^1 + s n_1^0}{n_2^1 + s n_2^0} - \frac{1}{2} \frac{n_1^1 + s n_1^0}{n_2^1 + s n_2^0} \sqrt{1 - 4 \left(1 - o\left(\frac{1}{s}\right) \right) \frac{n_2^1 + s n_2^0}{(n_1^1 + s n_1^0)^2}}.$$

The coefficient s is defined as $s = \frac{|n^0|}{|n^1|}$. Now we have to find the Taylor expansion of e^{β_0} of degree two in $1/s$ and reinject it in (B.1). We have

$$\begin{aligned} \frac{n_2^1 + s n_2^0}{(n_1^1 + s n_1^0)^2} &= \frac{n_2^1 + s n_2^0}{s^2 (n_1^0)^2} \frac{1}{\left(\frac{n_1^1}{s n_1^0} + 1\right)^2} = \frac{n_2^1 + s n_2^0}{s^2 (n_1^0)^2} \left(1 - 2 \frac{n_1^1}{s n_1^0} + o\left(\frac{1}{s}\right) \right) \\ &= \frac{1}{s} \frac{n_2^0}{(n_1^0)^2} + \frac{1}{s^2} \left(\frac{n_2^1}{(n_1^0)^2} - 2 \frac{n_1^1 n_2^0}{(n_1^0)^3} \right) + o\left(\frac{1}{s^2}\right), \end{aligned}$$

and

$$\begin{aligned} &\sqrt{1 - 4 \left(1 - o\left(\frac{1}{s}\right) \right) \left(\frac{1}{s} \frac{n_2^0}{(n_1^0)^2} + \frac{1}{s^2} \left(\frac{n_2^1}{(n_1^0)^2} - 2 \frac{n_1^1 n_2^0}{(n_1^0)^3} \right) + o\left(\frac{1}{s^2}\right) \right)} \\ &= \sqrt{1 - 4 \left(\frac{1}{s} \frac{n_2^0}{(n_1^0)^2} + \frac{1}{s^2} \left(\frac{n_2^1}{(n_1^0)^2} - 2 \frac{n_1^1 n_2^0}{(n_1^0)^3} \right) + o\left(\frac{1}{s^2}\right) \right)} \\ &= 1 - 2 \left(\frac{1}{s} \frac{n_2^0}{(n_1^0)^2} + \frac{1}{s^2} \left(\frac{n_2^1}{(n_1^0)^2} - 2 \frac{n_1^1 n_2^0}{(n_1^0)^3} + \frac{1(n_2^0)^2}{(n_1^0)^4} \right) \right) + o\left(\frac{1}{s^2}\right). \end{aligned}$$

Therefore

$$\begin{aligned} e^{\beta_0} &= \frac{n_1^1 + s n_1^0}{n_2^1 + s n_2^0} \left(\frac{1}{s} \frac{n_2^0}{(n_1^0)^2} + \frac{1}{s^2} \left(\frac{n_2^1}{(n_1^0)^2} - 2 \frac{n_1^1 n_2^0}{(n_1^0)^3} + \frac{1(n_2^0)^2}{(n_1^0)^4} \right) + o\left(\frac{1}{s^2}\right) \right) \\ &= \frac{n_1^1 + s n_1^0}{s n_2^0} \left(1 - \frac{n_2^1}{s n_2^0} + o\left(\frac{1}{s}\right) \right) \left(\frac{1}{s} \frac{n_2^0}{(n_1^0)^2} + \frac{1}{s^2} \left(\frac{n_2^1}{(n_1^0)^2} - 2 \frac{n_1^1 n_2^0}{(n_1^0)^3} + \frac{(n_2^0)^2}{(n_1^0)^4} \right) + o\left(\frac{1}{s^2}\right) \right) \\ &= \left(\frac{n_1^0}{n_2^0} + \frac{1}{s} \left(\frac{n_1^1}{n_2^0} - \frac{n_1^1 n_1^0}{(n_2^0)^2} \right) \right) \left(\frac{1}{s} \frac{n_2^0}{(n_1^0)^2} + \frac{1}{s^2} \left(\frac{n_2^1}{(n_1^0)^2} - 2 \frac{n_1^1 n_2^0}{(n_1^0)^3} + \frac{(n_2^0)^2}{(n_1^0)^4} \right) \right) + o\left(\frac{1}{s^2}\right) \end{aligned}$$

$$e^{\beta_0} = \frac{1}{s n_1^0} + \left(\frac{1}{s n_1^0} \right)^2 \left(\frac{n_2^0}{n_1^0} - n_1^1 \right) + o\left(\frac{1}{s^2}\right).$$

The system of equations (B.1) without its first equation is

$$\overline{N}^1 = \frac{\mathbf{I}_1^T n^1}{|n^1|} = e^{\beta_0} (\mathbf{n}_1^1 + s \mathbf{n}_1^0) - e^{2\beta_0} (\mathbf{n}_2^1 + s \mathbf{n}_2^0) + o\left(\frac{1}{s}\right)$$

and we use the previous expression for e^{β_0} :

$$\overline{N}^1 = \left(\frac{1}{s n_1^0} + \left(\frac{1}{s n_1^0} \right)^2 \left(\frac{n_2^0}{n_1^0} - n_1^1 \right) \right) (\mathbf{n}_1^1 + s \mathbf{n}_1^0) - \left(\frac{1}{s n_1^0} \right)^2 (\mathbf{n}_2^1 + s \mathbf{n}_2^0) + o\left(\frac{1}{s}\right)$$

$$\overline{N}^1 = \frac{\mathbf{n}_1^0}{n_1^0} + \frac{1}{s} \left(\frac{n_2^0}{(n_1^0)^2} \left[\frac{\mathbf{n}_1^0}{n_1^0} - \frac{\mathbf{n}_2^0}{n_2^0} \right] + \frac{n_1^1}{n_1^0} \left[\frac{\mathbf{n}_1^1}{n_1^1} - \frac{\mathbf{n}_1^0}{n_1^0} \right] \right) + o\left(\frac{1}{s}\right)$$

or

$$\overline{N}^1 + \frac{1}{s} \left(\frac{\mathbf{n}_2^0}{(n_1^0)^2} - \frac{\mathbf{n}_1^1}{n_1^0} \right) = \frac{\mathbf{n}_1^0}{n_1^0} \left(1 + \frac{1}{s} \left(\frac{n_2^0}{(n_1^0)^2} - \frac{n_1^1}{n_1^0} \right) \right) + o\left(\frac{1}{s}\right).$$

Then, using again a Taylor expansion,

$$\frac{\mathbf{n}_1^0}{n_1^0} = \overline{N}^1 + \frac{1}{s} \left(\frac{\mathbf{n}_2^0 - \overline{N}^1 n_2^0}{(n_1^0)^2} - \frac{\mathbf{n}_1^1 - \overline{N}^1 n_1^1}{n_1^0} \right) + o\left(\frac{1}{s}\right)$$

or

$$\frac{\mathbf{n}_1^0}{n_1^0} - \overline{N}^1 = \frac{1}{s} \left(\frac{n_2^0}{(n_1^0)^2} \left(\frac{\mathbf{n}_2^0}{n_2^0} - \overline{N}^1 \right) - \frac{n_1^1}{n_1^0} \left(\frac{\mathbf{n}_1^1}{n_1^1} - \overline{N}^1 \right) \right) + o\left(\frac{1}{s}\right).$$

C Path simulations

In the following graphs, the dotted lines are obtained by the exact path corresponding to theorem 4.1 for examples 1 and 2, theorem 4.2 for examples 3 and 4 and the inclusion case for examples 5 and 6. The solid lines are always given by a coordinate descent algorithm for the standard logistic regression. We change the scale for t (by a linear rescaling) in order to have to same $\max(t)$ (see (4.8)) for the exact and algorithmic paths.




Figure 3: Example 1. Independence and class imbalance. $\beta_0 = -3$ ($|n^0|/|n^1| = 17$).




Figure 4: Example 2. Independence and no class imbalance. $\beta_0 = 0$ ($|n^0|/|n^1| = 1.5$).




Figure 5: Example 3. Orthogonality and class imbalance. $\beta_0 = -3$ ($|n^0|/|n^1| = 15$).




Figure 6: Example 4. Orthogonality and no class imbalance. $\beta_0 = 0$ ($|n^0|/|n^1| = 0.8$).




Figure 7: Example 5. Inclusion and class imbalance. $\beta_0 = -3$ ($|n^0|/|n^1| = 15$).




Figure 8: Example 6. Inclusion and no class imbalance. $\beta_0 = 0$ ($|n^0|/|n^1| = 0.8$).

References

- [1] AHMED ISMAIL, PARIENTE ANTOINE, TUBERT-BITTER PASCALE (2016) Class-imbalanced subsampling lasso algorithm for discovering adverse drug reactions. *Statistical Methods in Medical Research*.
- [2] BEZIZ ET AL. (2016) Spontaneous adverse drug reaction reporting in France: A retrospective analysis of reports made to the French medicines agency from 2002 to 2014. *Revue d'Épidémiologie et de Santé Publique*, **64**.
- [3] CASTER ET AL. (2010) Large-Scale Regression-Based Pattern Discovery: The Example of Screening the WHO Global Drug Safety Database. *Stat. Anal. Data Min.*, **3**, no. 4, 197–208.
- [4] COX DAVID ROXBEE (1975) Partial likelihood. *Biometrika*, **62**, no. 2, 269–276.
- [5] EFRON BRADLEY, HASTIE TREVOR, JOHNSTONE IAIN, TIBSHIRANI ROBERT (2004) Least angle regression. *Annals of statistics*, **32**, no. 2, 407–499.
- [6] ELRAHMAN SHAZA, ABRAHAM AJITH (2013) A Review of Class Imbalance Problem. *Journal of Network and Innovative Computing*, **1**, 332–340.
- [7] FIRTH DAVID (1993) Bias reduction of maximum likelihood estimates. *Biometrika*, **80**, no. 1, 27–38.
- [8] FITHIAN WILLIAM, HASTIE TREVOR. (2014) Local case-control sampling: efficient subsampling in imbalanced data sets. *Ann. Statist.*, **42**, no. 5, 1693–1724.
- [9] FRIEDMAN JEROME, HASTIE TREVOR, TIBSHIRANI ROB (2010) Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, **33**, no. 1, 1–22.
- [10] FRIEDMAN JEROME, HASTIE TREVOR, HFLING HOLGER, TIBSHIRANI ROBERT (2007) Regularization Paths for Generalized Linear Models via Coordinate Descent. *Annals of Applied Statistics*, **1**, no. 2, 302–332.
- [11] GOVAERT GÉRARD, NADIF MOHAMED (2008) Block clustering with Bernoulli mixture models: comparison of different approaches. *Comput. Statist. Data Anal.*, **52**, no. 6, 3233–3245.
- [12] GUO XINJIAN ET AL. (2008) On the Class Imbalance Problem. *Proceedings of the Fourth International Conference on Natural Computation*. **4**, 192–201.
- [13] HARPAZ ET AL. (2013) Performance of pharmacovigilance signal-detection algorithms for the FDA adverse event reporting system. *Clin Pharmacol. Ther.*, **6**
- [14] HIRIART-URRUTY JEAN-BAPTISTE (1981) A note on the Legendre-Fenchel transform of convex composite functions. *Nonsmooth Mechanics and Analysis. Adv. Mech. Math. Springer, New York*, **12**, 35–46.
- [15] JEFFREYS HAROLD. (1946) An invariant form for the prior probability in estimation problems. *Proc. Roy. Soc. London. Ser. A.*, **186**, 453–461.
- [16] KI HANG KIM (1982) Boolean matrix theory and applications. *Monographs and textbooks in pure and applied mathematics*. ISBN-13: 978-0824717889

- [17] LIPOVETSKY STAN (2015) Analytical closed-form solution for binary logit regression by categorical predictors. *J. Appl. Stat.*, **42**, no. 1, 37–49.
- [18] MADIGAN DAVID, RYAN PATRICK, SIMPSON SHAWN, ZORYCH IVAN (2011) Bayesian methods in pharmacovigilance. With discussion by William DuMouchel. *Oxford Univ. Press*, Bayesian statistics 9, 421–438.
- [19] MANTEL, N. AND HAENSZEL, W. (1959) Statistical aspects of the analysis of data from retrospective studies of disease. *J. Natl. Cancer Inst.*, **22**, 719–748.
- [20] MINKA THOMAS P. (2007) A comparison of numerical optimizers for logistic regression. URL: <http://research.microsoft.com/en-us/um/people/minka/papers/logreg/>
- [21] OOMMEN THOMAS ET AL. (2011) Sampling Bias and Class Imbalance in Maximum-likelihood Logistic Regression. *Mathematical Geosciences.*, **43**, no. 1, 99–120.
- [22] OWEN ART B. (2007) Infinitely Imbalanced Logistic Regression. *J. Mach. Learn. Res.*, **8**, 761–773.
- [23] ROSSET SAHARON (2004). Following Curved Regularized Optimization Solution Paths. *Advances in Neural Information Processing Systems*. **17**, 1153–1160.
- [24] ROSSET S. AND ZHU J. (2007). Piecewise linear regularized solution paths. *The Annals of Statistics*. **35**, no. 3, 1012–1030.
- [25] SAINT-RAYMOND LAURE (2009). Hydrodynamic Limits of the Boltzmann Equation. *Springer. Lecture Notes in Mathematics*. ISBN: 978-3-540-92846-1
- [26] SCHWARZ GIDEON (1978) Estimating the dimension of a model. *The Annals of Statistics*, **6**, no. 2, 461–464.
- [27] SILVAPULLE MERVYN J. (1981) On the Existence of Maximum Likelihood Estimators for the Binomial Response Models. *J. Roy. Statist. Soc. Ser. B*, **43**, no. 3, 310–313.
- [28] SIMPSON SHAWN E, MADIGAN DAVID, ZORYCH IVAN, SCHUEMIE MARTIJN J., RYAN PATRICK B., SUCHARD MARC A. (2013) Multiple self-controlled case series for large-scale longitudinal observational databases. *Biometrics*, **69**, no. 4, 893–902.
- [29] TIBSHIRANI ROBERT (1996) Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, **58**, no. 1, 267–288.
- [30] YU HSIANG-FU, HUANG FANG-LAN, LIN CHIH-JEN (1981) Dual coordinate descent methods for logistic regression and maximum entropy models. *Mach. Learn.*, **85**, no. 1-2, 41–75.
- [31] ZOU HUI, HASTIE TREVOR (2005) Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **67**, no. 2, 301–320.